

IBM MQ



Installing IBM MQ

Version 8 Release 0

Note

Before using this information and the product it supports, read the information in "Notices" on page 1329.

This edition applies to version 8 release 0 of WebSphere MQ and to all subsequent releases and modifications until otherwise indicated in new editions.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright IBM Corporation 2007, 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables xi

Planning 1

Architectures based on a single queue manager.	1
Single queue manager with local applications accessing a service	2
Single queue manager with remote applications accessing a service as clients	2
Single queue manager with a publish/subscribe configuration	2
Architectures based on multiple queue managers	2
Planning your distributed queues and clusters	3
Planning your distributed publish/subscribe network	58
Planning your storage and performance requirements	98
Disk space requirements on distributed platforms	98
Planning file system support	100
IBM MQ and UNIX System V IPC resources	128
Shared memory on AIX	128
IBM MQ and UNIX Process Priority.	129
Planning your IBM MQ client environment on HP Integrity NonStop Server	129
Preparing the HP Integrity NonStop Server environment.	129
IBM MQ and HP NonStop TMF	130
Using HP NonStop TMF	130
Planning your IBM MQ environment on z/OS	132
Planning your storage and performance requirements on z/OS	132
Planning your page sets and buffer pools	138
Planning your coupling facility and offload storage environment	146
Planning your logging environment	157
Planning for IBM MQ Managed File Transfer	164
Planning for channel initiator SMF data	167
Planning for backup and recovery	169
Planning your z/OS UNIX or UNIX System Services environment.	178
Planning your z/OS TCP/IP environment.	179

Installing and uninstalling 181

Planning your installation	181
Choosing an installation name.	182
Multiple installations	183
Choosing a primary installation	184
Uninstalling, upgrading, and maintaining the primary installation	191
Choosing an installation location	192
Choosing what to install.	194
Installing IBM MQ Telemetry	214
Planning your installation on Windows systems	215

Planning your installation on HP Integrity NonStop Server	220
Checking requirements	224
Preparing the system	233
Operating System configuration and tuning for IBM MQ on AIX systems	239
Operating System configuration and tuning for IBM MQ on HP-UX systems	239
Operating system configuration and tuning for IBM MQ on Linux systems	242
Operating System configuration and tuning for IBM MQ on Solaris systems	245
Operating System configuration and tuning for IBM MQ on IBM i	247

Installing IBM MQ 249

Installation using Electronic Software Download	249
Redistributable clients	250
Installation considerations for redistributable clients	251
.NET application runtime - Windows only.	253
Installing an IBM MQ server	254
Installing IBM MQ server on AIX.	256
Installing IBM MQ server on HP-UX	259
Installing IBM MQ server on Linux	263
Installing IBM MQ server on Solaris.	272
Installing IBM MQ server on Windows.	277
Installing IBM MQ server on IBM i	303
Converting a trial license on UNIX, Linux, and Windows.	316
Displaying messages in your national language on UNIX and Linux systems	317
Displaying messages in your national language on Windows systems.	318
Installing an IBM MQ client	319
Installing an IBM MQ client on AIX systems	319
Installing an IBM MQ client on HP Integrity NonStop Server systems.	321
Installing an IBM MQ client on HP-UX systems	323
Installing an IBM MQ client on Linux	325

Installing an IBM MQ client on Solaris 333

Non-interactive installation of IBM MQ client on Solaris.	334
---	-----

Installing an IBM MQ client on Windows systems 337

Advanced installation using msixec	338
Specifying command line parameters with msixec	339
Using a response file with msixec	340
Multiple installation using MSI Instance ID	342
Using transforms with msixec	342
Creating a response file	343
Using the MQParms command	343

Modifying the client installation on Windows . . .	347
Modifying the client installation on Windows using Add/Remove Programs	347
Silently modifying an IBM MQ client installation using msisexec	348
Silently modifying an IBM MQ client installation using MQParms	348

Installing an IBM MQ client on IBM i 349

Installation of IBM MQ client and IBM MQ server for IBM i	351
---	-----

Installing IBM MQ Advanced Message Security 353

Installing IBM MQ Advanced Message Security on AIX	353
Installing using SMIT.	353
Installing using command line.	354
Installing IBM MQ Advanced Message Security on HP-UX	354
Installing IBM MQ Advanced Message Security on Linux	354
Installing IBM MQ Advanced Message Security on Windows	355
Using the Launchpad.	355
Installing IBM MQ Advanced Message Security on IBM i	356
Installing IBM MQ Advanced Message Security on z/OS	357

Installing IBM MQ Java messaging and web services for IBM i 359

Verifying an IBM MQ installation 363

Verifying a server installation	363
Verify the installation using the command line	364
Verify the installation using the Postcard application	370
Verifying a client installation	374
Verifying a client installation using the command line	374
Verifying a client installation using IBM MQ Explorer	379
Testing communication between a client and a server	382
Verifying the installation of IBM MQ Telemetry	384
Verifying the installation of IBM MQ Telemetry by using MQ Explorer	385
Verifying the installation of IBM MQ Telemetry using the command line.	387

Uninstalling IBM MQ components 389

Uninstalling	389
Uninstalling IBM MQ on AIX	389
Uninstalling IBM MQ on HP Integrity NonStop Server	391
Uninstalling IBM MQ on HP-UX	391
Uninstalling IBM MQ on Linux	392
Uninstalling IBM MQ on Solaris	394

Uninstalling IBM MQ on Windows systems . . .	395
Uninstalling IBM MQ for IBM i	400
Uninstalling IBM MQ Telemetry components. . .	405
Uninstalling IBM MQ Advanced Message Security	405
Uninstalling on AIX	405
Uninstalling on HP-UX	407
Uninstalling on Linux	407
Uninstalling on Windows	408

Installing IBM MQ for z/OS. 409

Planning to install IBM MQ	412
Delivery media.	413
Customizing IBM MQ and its adapters.	413
Verifying your installation of IBM MQ for z/OS	413
Macros intended for customer use	413
Sub-capacity license charges with IBM MQ for z/OS	414
IBM MQ for z/OS Value Unit Edition (VUE). . .	415
IBM MQ Advanced Message Security for z/OS .	416
IBM MQ Managed File Transfer for z/OS	417

Migrating and upgrading IBM MQ. 419

Introduction to IBM MQ migration	421
Migration paths	421
The version naming scheme for IBM MQ for z/OS	422
The version naming scheme for IBM MQ (On platforms other than z/OS)	423
Internet Protocol Version 6 (IPv6) migration .	425
Overview of migration methods	434
Maintenance, upgrade, and migration	440
Coexistence, compatibility, and interoperability	456
Queue manager migration	472
Reverting a queue manager to a previous version	473
IBM MQ MQI client migration	474
Application migration and interoperation . .	477
Queue manager cluster migration	478
Queue-sharing group migration	480
Migrating a queue manager in a high-availability configuration.	481
Introduction to changes for Windows on IBM MQ Version 8.0	483
IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i	485
z/OS: Migration planning to the latest release .	486
z/OS: Review and modify queue manager customizations from the previous release . .	489
z/OS: IBM MQ Version 8.0 JCL changes	494
UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version .	496
UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version	500
UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version .	505
Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0	510
Windows: Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0	510

Linux: Migrating from IBM MQ Telemetry Version 7.0.1 to Version 8.0	511	Migrating from a single instance to a multi-instance queue manager	575
Migrating a queue manager to the latest release	512	Reverting to a single-instance queue manager	579
UNIX systems - migrating a queue manager from your current version to the latest version	513	Changes that affect migration	580
Windows: Migrating a queue manager from a previous version to the latest version	515	Coexistence	581
IBM i Migrating a queue manager from a previous release	517	Changes from WebSphere MQ Version 7.0.1 to IBM MQ Version 8.0	584
IBM i: migrating a queue manager from a previous release - alternative method	528	IBM MQ maintenance tasks (On platforms other than z/OS)	623
Restoring a queue manager from the latest version to a previous version on UNIX systems and Windows	531	Applying and removing maintenance level updates (On platforms other than z/OS)	623
Linux: Cleaning up after using the rpm freshen or upgrade options	533	Applying maintenance level upgrades to multi-instance queue managers	653
Migrating IBM MQ library loading from an earlier version of the product to the latest version	533	Migrating queue managers to new-function fix packs	655
Migrating an IBM MQ MQI client to the latest version of the product	535	Querying the maintenance level	669
Migrating an IBM MQ MQI client on UNIX systems, Windows, and IBM i to the latest version	536	UNIX, Linux, and Windows: Staging maintenance fixes	669
Restoring an IBM MQ MQI client and client connection to the previous version	537	Migration commands, utilities, and reference information	673
Migrating applications to the latest version of the product	537	JMS PROVIDERVERSION property	673
Migrating IBM MQ library loading from an earlier version of the product to the latest version	538	strmqbrk : Migrate the IBM WebSphere MQ Version 6.0 publish/subscribe broker to a later version	676
Linux: Rebuilding a C++ application	545	z/OS: OPMODE	677
Migrating IBM MQ for z/OS - order of tasks	547	z/OS: Switching from OPMODE=(NEWFUNC,800) to OPMODE=(COMPAT,800)	679
z/OS: new messages in IBM MQ for z/OS Version 8	548	PROPCTL channel options	680
z/OS: overall migration - order of tasks	550	PROPCTL queue options	682
Migrating from earlier unsupported releases of IBM MQ for z/OS	551	MQGM0 message property option settings	685
Backward migration to earlier supported releases of IBM MQ for z/OS	552	Configuring	687
Preparing to migrate a single IBM MQ for z/OS queue manager	554	Creating and managing queue managers on distributed platforms	687
Migrating a single IBM MQ z/OS queue manager to the next release of the product	558	Creating a default queue manager	690
Post migration tasks	564	Making an existing queue manager the default	692
z/OS: Adding a new queue-sharing group to an existing Db2 data sharing group in the latest version	566	Backing up configuration files after creating a queue manager	692
z/OS: Migrating queue sharing groups from a previous version of the product	567	Starting a queue manager	693
Migrating a queue manager cluster	569	Stopping a queue manager	694
Migrating a queue manager cluster: Create a plan	569	Restarting a queue manager	695
Migrating a queue manager cluster: Create a backout plan	570	Deleting a queue manager	695
Migrating a queue manager cluster: Migrating one cluster queue manager	571	Configuring connections between the server and client	696
Migrating a queue manager cluster: Migrating the test system	572	Which communication type to use	698
Migrating a queue manager cluster: Migrating the production system	573	Configuring an extended transactional client	701
Windows: Migrating an MSCS configuration	573	Defining MQI channels	712
		Creating server-connection and client-connection definitions on different platforms	713
		Creating server-connection and client-connection definitions on the server	716
		Channel-exit programs for MQI channels	721
		Connecting a client to a queue-sharing group	726
		Configuring a client using a configuration file	727
		Using IBM MQ environment variables	747
		Changing IBM MQ and queue manager configuration information	755
		Changing configuration information on Windows, UNIX and Linux systems	756
		Changing configuration information on IBM i	763

Attributes for changing IBM MQ configuration information	775
Changing queue manager configuration information	782
Configuring distributed queuing	802
IBM MQ distributed queuing techniques	802
Introduction to distributed queue management	823
Monitoring and controlling channels on Windows, UNIX and Linux platforms	856
Monitoring and controlling channels on IBM i	879
Configuring a queue manager cluster	902
Configuring publish/subscribe messaging	1004
Setting queued publish/subscribe message attributes	1004
Starting queued publish/subscribe	1005
Stopping queued publish/subscribe	1006
Adding a stream	1006
Deleting a stream	1007
Adding a subscription point	1008
Configuring distributed publish/subscribe networks	1009
Configuring multiple installations	1023
Connecting applications in a multiple installation environment	1024
Changing the primary installation	1033
Associating a queue manager with an installation	1034
Finding installations of IBM MQ on a system	1036
Availability, recovery and restart	1037
Automatic client reconnection	1038
Console message monitoring	1044
Using IBM MQ with high availability configurations	1047
Making sure that messages are not lost (logging)	1130
Backing up and restoring IBM MQ queue manager data	1150
Configuring JMS resources	1155
Configuring connection factories and destinations in a JNDI namespace	1156

Configuring JMS objects using MQ Explorer	1158
Configuring JMS objects using the administration tool	1159
Configuring JMS resources in WebSphere Application Server	1167
Configuring the application server to use the latest resource adapter maintenance level.	1174
Configuring the JMS PROVIDERVERSION property	1176
Configuring HP Integrity NonStop Server	1186
Gateway process overview	1187
Configuring Gateway to run under Pathway	1187
Configuring the client initialization file	1189
Granting permissions to channels	1189
Configuring IBM MQ using Docker	1189
Docker support on Linux systems	1190
Planning your own IBM MQ queue manager image using Docker	1190
Building a sample IBM MQ queue manager image using Docker	1191
Configuring queue managers on z/OS	1194
Preparing to customize your IBM MQ for z/OS queue managers	1194
Customizing IBM MQ for z/OS	1198
Testing your queue manager on z/OS	1252
Setting up communications with other queue managers	1261
Using IBM MQ with IMS	1294
Using IBM MQ with CICS.	1303
Upgrading and applying service to Language Environment or z/OS Callable Services	1303
Using OTMA exits in IMS.	1305

Index 1309

Notices 1329

Programming interface information	1330
Trademarks	1331

Sending your comments to IBM 1333

Figures

1. Network diagram showing all channels	5	33. Propagation of publications through a queue manager network	95
2. Network diagram showing QM-concentrators	7	34. On IBM MQ server 1	103
3. Distributed queuing.	14	35. On IBM MQ server 1	103
4. A network of four queue managers	15	36. On IBM MQ server 2	103
5. Two connected full repositories.	16	37. On IBM MQ server 1	104
6. More than two connected full repositories	18	38. On IBM MQ server 2	104
7. Classes of service	21	39. On IBM MQ server 1	104
8. A small cluster of two queue managers	24	40. On IBM MQ server 2	106
9. A cluster of queue managers	25	41. On IBM MQ server 2	106
10. A cluster of queue managers with sender channels	26	42. Output from a successful run of amqsfhac	110
11. A cluster of queue managers, showing auto-defined channels	27	43. Overall view of IBM MQ directory structure	111
12. Client-server application deployed to hub and spoke architecture using IBM MQ clusters	36	44. Example IPC subdirectory, pre-V7.0.1	112
13. Example of specific transmission queues for different departmental IBM MQ clusters	39	45. Example IPC subdirectory, V7.0.1 and subsequent releases	112
14. Example of specific transmission queues for different departmental IBM MQ clusters	43	46. Example: setting MQC_IPC_HOST	112
15. Definitions for the basic clusters	46	47. Example default IBM MQ directory structure for UNIX and Linux systems	113
16. Changes to isolate the sales queue in a new cluster and separate the gateway cluster transmission queues	47	48. Default directory structure (UNIX systems) after a queue manager has been started	115
17. Remove the sales queue on queue manager SALESRV from the sales cluster.	47	49. Directory structure pattern template	118
18. A queue manager cluster	64	50. Example default IBM MQ directory structure for UNIX and Linux systems	119
19. A direct routed publish/subscribe cluster	65	51. Typical directory structure for releases earlier than v7.0.1	120
20. A direct routed publish/subscribe cluster with a publisher and a subscriber to a clustered topic	66	52. Share qmgrs and log directories	120
21. A direct routed publish/subscribe cluster that is fully interconnected	67	53. Share named qmgrs and log directories	121
22. A queue manager cluster	69	54. Share everything	122
23. A topic host routed publish/subscribe cluster with one topic defined on one topic host.	70	55. How IBM MQ stores long messages on page sets	141
24. A topic host routed publish/subscribe cluster with one topic defined on one topic host, and one subscriber	71	56. Calculating the size of a coupling facility structure	149
25. A topic host routed publish/subscribe cluster with one topic, one subscriber and one publisher	72	57. Example of queue manager backup activity	177
26. A topic host routed publish/subscribe cluster with two topics, each defined on one topic host	73	58. IBM MQ application migration model	450
27. Creating proxy subscriptions in a multiple topic host publish/subscribe cluster	74	59. Side-by-side installation - step 2	457
28. Receiving publications in a multiple topic host publish/subscribe cluster	75	60. Side-by-side installation - step 4	458
29. Routing publications to subscribers in a multiple topic host publish/subscribe cluster	76	61. Rolling fix packs	459
30. Overlapping clusters: Two clusters each subscribing to different topics	90	62. Coexistence of two queue managers using Version 7.0.1 and a later version installations	460
31. Propagation of subscriptions through a queue manager network	94	63. Loading calls in a different library	462
32. Multiple subscriptions	94	64. Running setmqenv	464
		65. Add Put1DefaultAlwaysSync to mqclient.ini	472
		66. Linux C server application, 32 bit, threaded compile and link Version 7.0.1	543
		67. Set TolerateRepositoryFailure to TRUE in qm.ini	596
		68. Default dspmqver options in IBM WebSphere MQ Version 7.0.1	600
		69. Default dspmqver options in IBM WebSphere MQ Version 7.1	600
		70. dspmqver with option to make IBM WebSphere MQ Version 7.1, or later, similar to IBM WebSphere MQ Version 7.0.1.	600
		71. Migration of queue managers to new command levels using new-function fix packs	656

72. Initial state, QM1 and QM2 at command level 7r0, and fix level 7.r.0.0	659	113. First of the set of channel status panels	886
73. QM1 and QM2 at command level 7r0, and fix level 7.r.0.1.	661	114. LU 6.2 communication setup panel - initiating end	897
74. QM1 at command level 7r0 and fix level 7.r.0.2 ; QM2 at command level 7r1 and fix level 7.r.0.2.	663	115. LU 6.2 communication setup panel - initiated end	900
75. QM1 at command level 7r0 and fix level 7.r.0.3 ; QM2 at command level 7r5 and fix level 7.r.0.3.	665	116. LU 6.2 communication setup panel - initiated end	901
76. QM1 at command level 7r0 and fix level 7.r.0.3 ; QM2 at command level 7R0 and fix level 7.R.0.0.	667	117. Transmission queue / cluster-sender channel precedence	909
77. Rolling fix releases	670	118. Multiple cluster sender channels	912
78. Simple channel definition	715	119. The INVENTORY cluster with two queue managers	914
79. Defining the server-connection channel	719	120. The INVENTORY cluster with three queue managers	925
80. Defining the client-connection channel	720	121. The INVENTORY cluster with three queue managers	927
81. Client connection-initiated exchange with agreement for client connection using security parameters	724	122. The INVENTORY cluster with three queue managers	930
82. Example of an IBM MQ configuration file for UNIX systems	759	123. The INVENTORY cluster with four queue managers	932
83. Example queue manager configuration file for IBM MQ for UNIX and Linux systems	761	124. Cluster and queue-sharing group	934
84. Example of an IBM MQ configuration file	773	125. The INVENTORY cluster with the full repository moved to PARIS	937
85. Example queue manager configuration file	774	126. A hub and spoke network	940
86. A remote queue definition is used to resolve a queue name to a transmission queue to an adjacent queue manager	806	127. Interconnected clusters	945
87. The remote queue definition allows a different transmission queue to be used	807	128. Client-server application deployed to hub and spoke architecture using IBM MQ clusters	948
88. Receiving messages directly, and resolving alias queue manager name	808	129. Using a queue manager alias to return the reply message to a different cluster	950
89. Three methods of passing messages through your system	809	130. Client-server application deployed to hub and spoke cluster architecture using remote queue definitions	953
90. Separating messages flows	811	131. Client-server application deployed to hub and spoke architecture using an additional cluster transmission queue.	955
91. Combining message flows on to a channel	812	132. Using an additional cluster to separate message traffic in the gateway queue manager that goes to one of a number of cluster queues on the same queue manager	958
92. Diverting message streams to another destination	813	133. Client-server application deployed to hub and spoke architecture with separate cluster transmission queues on the gateway queue manager.	962
93. Reply-to queue name substitution during PUT call	814	134. Putting from a queue manager outside the cluster	976
94. Reply-to queue alias example	816	135. Putting from a queue manager outside the cluster	979
95. Distributed queue management model	825	136. Putting to a queue manager outside the cluster	980
96. Create a queue (1)	830	137. Putting from a queue manager outside the cluster	982
97. Create a queue (2)	830	138. Bridging across clusters	983
98. Create a queue (3)	831	139. Using a queue manager alias to return the reply message to a different cluster	986
99. Create a queue (4)	831	140. A cluster with multiple instances of the same queue	990
100. Channel states and substates	835	141. The INVENTORY cluster, with three queue managers	992
101. Flows between channel states	836	142. The INVENTORY cluster, with four queue managers	994
102. Flows between channel states	837		
103. What happens when a message cannot be delivered	847		
104. The concepts of triggering	849		
105. Create channel (1)	881		
106. Create channel (2)	881		
107. Create channel (3)	882		
108. Create channel (4)	882		
109. Work with channels	883		
110. Display a TCP/IP channel (1)	884		
111. Display a TCP/IP channel (2)	885		
112. Display a TCP/IP channel (3)	885		

143. The INVENTORY cluster, with four queue managers	996	165. PROVIDERVERSION normal mode	1179
144. The PRICECHECK cluster, with four server queue managers, two repositories, and two query queue managers	998	166. PROVIDERVERSION normal mode with restrictions	1180
145. The INVENTORY cluster with four queue managers	1000	167. PROVIDERVERSION migration mode	1181
146. Connecting clusters using hierarchies	1017	168. PROVIDERVERSION unspecified	1183
147. Bridged clusters	1019	169. PROVIDERVERSION unspecified (continued)	1184
148. Overlapping clusters, non-overlapping topic spaces	1021	170. Sample IEFSSNss statements for defining subsystems	1206
149. Connecting applications in a multiple installation environment	1025	171. RACF commands for CSQ4IVP1	1254
150. Automatic client reconnection.	1040	172. Sample report from CSQ4IVP1	1256
151. HA cluster	1050	173. RACF commands for CSQ4IVP1 for a queue-sharing group.	1257
152. Shared named data and log directories	1051	174. RACF commands for CSQ4IVPX.	1259
153. Two-computer MSCS cluster	1059	175. Example output from CSQ4IVPX	1260
154. Multi-instance queue manager	1075	176. The operations and controls initial panel	1267
155. Securing queue manager data and logs using an alternative global security group (1)	1104	177. Listing channels	1268
156. Securing queue manager data and logs using an alternative global security group (2)	1105	178. Command Entry	1271
157. Securing queue manager data and logs using an alternative global security group (3)	1105	179. Command Output	1271
158. Securing queue manager data and logs using an alternative global security principal (4)	1106	180. Starting a system function	1272
159. Securing queue manager data and logs using an alternative local security group, wmq.	1107	181. Stopping a function control	1273
160. Checkpointing	1135	182. Starting a channel.	1275
161. Checkpointing with a long-running transaction	1136	183. Testing a channel	1277
162. JMS objects created in IBM MQ	1157	184. Stopping a channel	1279
163. Objects created in WebSphere Application Server, and the corresponding objects in IBM MQ	1157	185. Listing channel connections	1281
164. Messaging provider modes	1177	186. Listing cluster channels.	1282
		187. Sample JCL to link-edit the dynamic call stub	1296
		188. CSQQDEFX macro syntax	1300
		189. Layout of a queue manager definition table	1301
		190. Example transaction definition for CSQQTRMN	1301
		191. Example PSB definition for CSQQTRMN	1301
		192. Example SMP/E LINK CALLLIBS job	1305
		193. OTMA pre-routing exit assembler sample	1307
		194. Sample assembler DRU exit	1309

Tables

1. Example of channel names	5	39. Features installed with each type of interactive installation.	216
2. Definitions for distributed queuing.	15	40. Required levels of Microsoft.NET	217
3. Definitions for clustering	16	41.	221
4. Queue manager names and port numbers	45	42.	223
5. Cluster sender and receiver channels for each routing method.	80	43.	224
6. Publish/subscribe system queues on distributed platforms	96	44. Minimum tunable kernel parameters values	240
7. Attributes of publish/subscribe system queues	96	45. Minimum tunable kernel parameters values	242
8. Running the data integrity check on two servers at the same time	104	46. Package component dependencies	270
9. Successful locking on two servers.	106	47. msiexec property=value parameters	281
10. Successful locking on two servers - verbose mode	108	48. Response file parameters.	282
11. Documented content of the /var/mqm directory on UNIX systems	116	49. Supplied transform files for various language support	286
12. Documented contents of the /var/mqm/qmgrs/qmname directory on UNIX systems	116	50. Properties used by MQParms in the MSI stanza	289
13. Documented contents of the /var/mqm/log/qmname directory on UNIX systems	117	51. Valid values for the MQPLANGUAGE property	290
14. Directories and files in the <i>FilePath</i> directory	124	52. Properties used in the Services stanza	291
15. Directories and files in <i>DefaultPrefix</i> directory	124	53. Globalizations of IBM MQ for IBM i.	304
16. Directories and files in <i>DataPath</i> and <i>Prefix/Qmgrs/QmgrName</i> directories	125	54. Globalizations of IBM MQ for IBM i.	308
17. Suggested definitions for JCL region sizes	137	55.	310
18. Where to find more information about storage requirements	138	56. Language identifiers	317
19. Buffer pool characteristics by OPMODE setting	144	57. AIX specific language identifiers	318
20. Suggested definitions for buffer pool settings	145	58. Package component dependencies	330
21. Minimum administrative structure sizes	148	59. msiexec property=value parameters	340
22. Calculating the size of a coupling facility structure	150	60. Response file parameters.	341
23. Table showing CSQSYSAPPL usage against sizing.	151	61. Supplied transform files for various language support	343
24. Planning your Db2 storage requirements	155	62. Properties used by MQParms in the MSI stanza	345
25.	166	63. Valid values for the MQPLANGUAGE property	346
26.	166	64. Globalizations of IBM MQ for IBM i.	401
27.	166	65. Migration paths: IBM MQ (On platforms except z/OS)	421
28. Primary installation options.	185	66. Migration paths: IBM MQ for z/OS	421
29. Installation location of IBM MQ	192	67. Effects of CONNAME and LOCLADDR settings	430
30. IBM MQ filesets for AIX systems	195	68. Abbreviations used in system configurations	432
31. IBM MQ message catalogs for AIX systems	197	69. System configurations.	432
32. IBM MQ components for HP-UX systems	198	70. Options for loading libraries	440
33. IBM MQ message catalogs for HP-UX systems	200	71. Types of upgrade (Platforms other than z/OS)	455
34. IBM MQ components for Linux systems	200	72. IBM MQ for z/OS additional members	494
35. IBM MQ message catalogs for Linux systems	203	73. IBM MQ for z/OS changed members	495
36. IBM MQ components for Solaris systems	204	74. National-language versions of IBM MQ for IBM i	524
37. IBM MQ message catalogs for Solaris systems	205	75. Windows configurations	540
38. IBM MQ interactive installation options for Solaris systems	206	76. UNIX and Linux configurations	543
		77. Before migration	547
		78. Migrating to the next release	547
		79. Post migration tasks	548
		80. Overview of migration	550
		81. Migrating your system to the next release	550
		82. Migrating a four-node MSCS cluster	574

83. Compatibility of queue manager versions with operations and control panel versions	583	123. Default outstanding connection requests (SPX)	796
84. JMS: API support	585	124. Three ways of using the remote queue definition object.	805
85. JMS: Client communication (not running on top of C client)	586	125. Reply-to queue alias	818
86. Java: API support	587	126. Queue name resolution at queue manager QMA	821
87. Java: Client communication (not running on top of C client)	588	127. Queue name resolution at queue manager QMB	821
88. XMS .NET: API support	589	128. Reply-to queue name translation at queue manager QMA	822
89. XMS .NET: Client communication (not running on top of C client)	590	129. Functions required in Windows, UNIX and Linux systems	857
90. XMS C and C++: API support	590	130. Settings on the local Windows system for a remote queue manager platform	868
91. XMS C and C++: Client communication (not running on top of C client)	591	131. Maximum outstanding connection requests queued at a TCP/IP port	875
92. .NET: API support	592	132. Settings on the local UNIX and Linux system for a remote queue manager platform	877
93. .NET: Client communication (not running on top of C client)	593	133. Settings on the local IBM i system for a remote queue manager platform	896
94. New fields added to existing data types	598	134. Parameter values for creating clusters 1 and 2	949
95. Renamed GSKit commands	603	135. Publish/subscribe configuration parameters	1004
96. IBM WebSphere MQ Version 7.1 or later JRE addressing modes	604	136. Benefits and drawbacks of the options for loading libraries	1027
97. Channel status fields affected by changes to subject and issuer distinguished names	612	137. Configuring applications to use particular assemblies	1031
98. Channel data structures affected by changes to subject and issuer distinguished names	612	138. Supported clients	1038
99. Default symbolic links created in releases before Version 7.1	618	139. Automatic reconnection configuration requirements	1039
100. Applying and removing maintenance	623	140. The JMS object types that are handled by the administration tool	1159
101. IBM MQ commands to display the installed versions	632	141. Administration verbs	1163
102. Properties used to install or uninstall a maintenance update	647	142. Syntax and description of commands used to manipulate subcontexts	1164
103. Properties used to install or uninstall a maintenance update	651	143. Syntax and description of commands used to manipulate administered objects	1167
104. Functions and capabilities in Version 7.0.1	679	144. IBM MQ libraries that exist after installation	1196
105. Functions and capabilities in Version 7.1	679	145. National language feature libraries	1199
106. Functions and capabilities in Version 8.0	679	146. Customization summary	1200
107. Channel message property attribute settings	681	147. Subsystem name to CPF associations	1207
108. Queue message property attribute settings	682	148. Example of CPF subset and superset rules	1207
109. MQGMO message property option settings	685	149. Valid character set for CPF strings	1208
110. Transmission protocols - combination of IBM MQ MQI client and server platforms.	697	150. Members of thlqual.SCSQPROC	1217
111. Transmission protocols - combination of IBM MQ client and server platforms	699	151. Members of thlqual.SCSQPROC	1218
112. Maximum outstanding connection requests queued at a TCP/IP port	700	152. Default values of CSQ6SYSP parameters	1223
113. Assumed values of the AXLIB parameter	707	153. Default values of CSQ6LOGP parameters	1232
114. IBM MQ libraries containing the XA switch structures	708	154. Default values of CSQ6ARVP parameters	1236
115. The switch load files	710	155. Required user access to class CSFSERV and CSFKEYS profiles.	1252
116. CICS task termination exits	711	156. Channel tasks	1265
117. Identifying API calls	725	157. Settings on the local z/OS system for a remote queue manager platform	1287
118. Which attributes apply to each type of client	729	158. SSM specifications options.	1299
119. Automatic reconnection depends on the values set in the application and in the mqclient.ini file	736	159. Service has been applied or the product has been upgraded to a new release	1303
120. List of possible ISO CCSIDs.	766	160. One of the products has been updated to a new release in a new SMP/E environment and libraries	1304
121. List of possible ISO CCSIDs.	776		
122. Default outstanding connection requests (TCP)	793		

Planning

When planning your IBM® MQ environment, consider the support that IBM MQ provides for single and multiple queue manager architectures, and for point-to-point and publish/subscribe messaging styles. Also plan your resource requirements, and your use of logging and backup facilities.

Before you plan your IBM MQ architecture, familiarize yourself with the basic IBM MQ concepts. See IBM MQ Technical overview.

IBM MQ architectures range from simple architectures using a single queue manager, to more complex networks of interconnected queue managers. Multiple queue managers are connected together using distributed queuing techniques. For more information about planning single queue manager and multiple queue manager architectures, see the following topics:

- “Architectures based on a single queue manager”
- “Architectures based on multiple queue managers” on page 2
 - “Planning your distributed queues and clusters” on page 3
 - “Planning your distributed publish/subscribe network” on page 58

z/OS On IBM MQ for z/OS® you can use shared queues and queue sharing groups to enable you to implement workload balancing, and your IBM MQ applications to be scalable and highly available. For information about shared queues and queue-sharing groups, see Shared queues and queue-sharing groups.

For information about planning for multiple installations, storage and performance requirements, and use of clients, see the other subtopics.

Related concepts:

z/OS “Planning your IBM MQ environment on z/OS” on page 132

When planning your IBM MQ environment, you must consider the resource requirements for data sets, page sets, DB2®, Coupling Facilities, and the need for logging, and backup facilities. Use this topic to plan the environment where IBM MQ runs.

Related information:

Checking requirements

Making sure that messages are not lost (logging)

Availability, recovery and restart

Architectures based on a single queue manager

The simplest IBM MQ architectures involve the configuration and use of a single queue manager.

Before you plan your IBM MQ architecture, familiarize yourself with the basic IBM MQ concepts. See IBM MQ Technical overview.

A number of possible architectures using a single queue manager are described in the following sections:

- “Single queue manager with local applications accessing a service” on page 2
- “Single queue manager with remote applications accessing a service as clients” on page 2
- “Single queue manager with a publish/subscribe configuration” on page 2

Single queue manager with local applications accessing a service

The first architecture based on a single queue manager is where the applications accessing a service are running on the same system as the applications providing the service. An IBM MQ queue manager provides asynchronous intercommunication between the applications requesting the service and the applications providing the service. This means that communication between the applications can continue even if one of the applications is offline for an extended period of time.

Single queue manager with remote applications accessing a service as clients

The second architecture based on a single queue manager has the applications running remotely from the applications providing the service. The remote applications are running on different systems to the services. The applications connect as clients to the single queue manager. This means that access to a service can be provided to multiple systems through a single queue manager.

A limitation of this architecture is that a network connection must be available for an application to operate. The interaction between the application and the queue manager over the network connection is synchronous.

Single queue manager with a publish/subscribe configuration

An alternative architecture using a single queue manager is to use a publish/subscribe configuration. In publish/subscribe messaging, you can decouple the provider of information from the consumers of that information. This differs from the point to point messaging styles in the previously described architectures, where the applications must know information about the target application, for example the queue name to put messages on. Using IBM MQ publish/subscribe the sending application publishes a message with a specified topic based on the subject of the information. IBM MQ handles the distribution of the message to applications that have registered an interest in that subject through a subscription. The receiving applications also do not need to know anything about the source of the messages to receive them. For more information, see [Publish/subscribe messaging and Example of a single queue manager publish/subscribe configuration](#).

Related concepts:

“Planning” on page 1

When planning your IBM MQ environment, consider the support that IBM MQ provides for single and multiple queue manager architectures, and for point-to-point and publish/subscribe messaging styles. Also plan your resource requirements, and your use of logging and backup facilities.

Related information:

[Introduction to IBM MQ](#)

[Creating and managing queue managers on distributed platforms](#)

Architectures based on multiple queue managers

You can use distributed message queuing techniques to create an IBM MQ architecture involving the configuration and use of multiple queue managers.

Before you plan your IBM MQ architecture, familiarize yourself with the basic IBM MQ concepts. See [IBM MQ Technical overview](#).

An IBM MQ architecture can be changed, without alteration to applications that provide services, by adding additional queue managers.

Applications can be hosted on the same machine as a queue manager, and then gain asynchronous communication with a service hosted on another queue manager on another system. Alternatively,

applications accessing a service can connect as clients to a queue manager that then provides asynchronous access to the service on another queue manager.

Routes that connect different queue managers and their queues are defined using distributed queuing techniques. The queue managers within the architecture are connected using channels. Channels are used to move messages automatically from one queue manager to another in one direction depending on the configuration of the queue managers.

For a high level overview of planning an IBM MQ network, see “Designing distributed queue manager networks” on page 4.

For information about how to plan channels for your IBM MQ architecture, see IBM MQ distributed queuing techniques.

Distributed queue management enables you to create and monitor the communication between queue managers. For more information about distributed queue management, see Introduction to distributed queue management.

Related concepts:

“Planning” on page 1

When planning your IBM MQ environment, consider the support that IBM MQ provides for single and multiple queue manager architectures, and for point-to-point and publish/subscribe messaging styles. Also plan your resource requirements, and your use of logging and backup facilities.

Related information:

Creating and managing queue managers on distributed platforms

Planning your distributed queues and clusters

You can manually connect queues hosted on distributed queue managers, or you can create a queue manager cluster and let the product connect the queue managers for you. To choose a suitable topology for your distributed messaging network, you need to consider your requirements for manual control, network size, frequency of change, availability and scalability.

Before you begin

This task assumes that you understand what distributed messaging networks are, and how they work. For a technical overview, see Distributed queuing and clusters.

About this task

To create a distributed messaging network, you can manually configure channels to connect queues hosted on different queue managers, or you can create a queue manager cluster. Clustering enables queue managers to communicate with each other without the need to set up extra channel definitions or remote queue definitions, simplifying their configuration and management.

To choose a suitable topology for your distributed publish/subscribe network, you need to consider the following broad questions:

- How much manual control do you need over the connections in your network?
- How big will your network be?
- How dynamic will it be?
- What are your availability and scalability requirements?

Procedure

- Consider how much manual control you need over the connections in your network.

If you only need a few connections, or if individual connections need to be very precisely defined, you should probably create the network manually.

If you need multiple queue managers that are logically related, and that need to share data and applications, you should consider grouping them together in a queue manager cluster.

- Estimate how big your network needs to be.
 1. Estimate how many queue managers you need. Bear in mind that queues can be hosted on more than one queue manager.
 2. If you are considering using a cluster, add two extra queue managers to act as full repositories.

For larger networks, manual configuration and maintenance of connections can be very time consuming, and you should consider using a cluster.

- Consider how dynamic the network activity will be.

Plan for busy queues to be hosted on performant queue managers.

If you expect queues to be frequently created and deleted, consider using a cluster.
- Consider your availability and scalability requirements.
 1. Decide whether you need to guarantee high availability of queue managers. If so, estimate how many queue managers this requirement applies to.
 2. Consider whether some of your queue managers are less capable than others.
 3. Consider whether the communication links to some of your queue managers are more fragile than to others.
 4. Consider hosting queues on multiple queue managers.

Manually configured networks and clusters can both be configured to be highly available and scalable. If you use a cluster, you need to define two extra queue managers as full repositories. Having two full repositories ensures that the cluster continues to operate if one of the full repositories becomes unavailable. Make sure that the full repository queue managers are robust, performant, and have good network connectivity. Do not plan to use the full repository queue managers for any other work.

- Based on these calculations, use the links provided to help you decide whether to manually configure connections between queue managers, or to use a cluster.

What to do next

You are now ready to configure your distributed messaging network.

Related information:

[Configuring distributed queuing](#)

[Configuring a queue manager cluster](#)

Designing distributed queue manager networks

IBM MQ sends and receives data between applications, and over networks using Queue Managers and Channels. Network planning involves defining requirements to create a framework for connecting these systems over a network.

Channels can be created between your system and any other system with which you need to have communications. Multi-hop channels can be created to connect to systems where you have no direct connections. The message channel connections described in the scenarios are shown as a network diagram in Figure 1 on page 5.

Channel and transmission queue names

Transmission queues can be given any name. But to avoid confusion, you can give them the same names as the destination queue manager names, or queue manager alias names, as appropriate. This associates the transmission queue with the route they use, giving a clear overview of parallel routes created through intermediate (multi-hopped) queue managers.

It is not so clear-cut for channel names. The channel names in Figure 1 for QM2, for example, must be different for incoming and outgoing channels. All channel names can still contain their transmission queue names, but they must be qualified to make them unique.

For example, at QM2, there is a QM3 channel coming from QM1, and a QM3 channel going to QM3. To make the names unique, the first one might be named 'QM3_from_QM1', and the second might be named 'QM3_from_QM2'. In this way, the channel names show the transmission queue name in the first part of the name. The direction and adjacent queue manager name are shown in the second part of the name.

A table of suggested channel names for Figure 1 is given in Table 1.

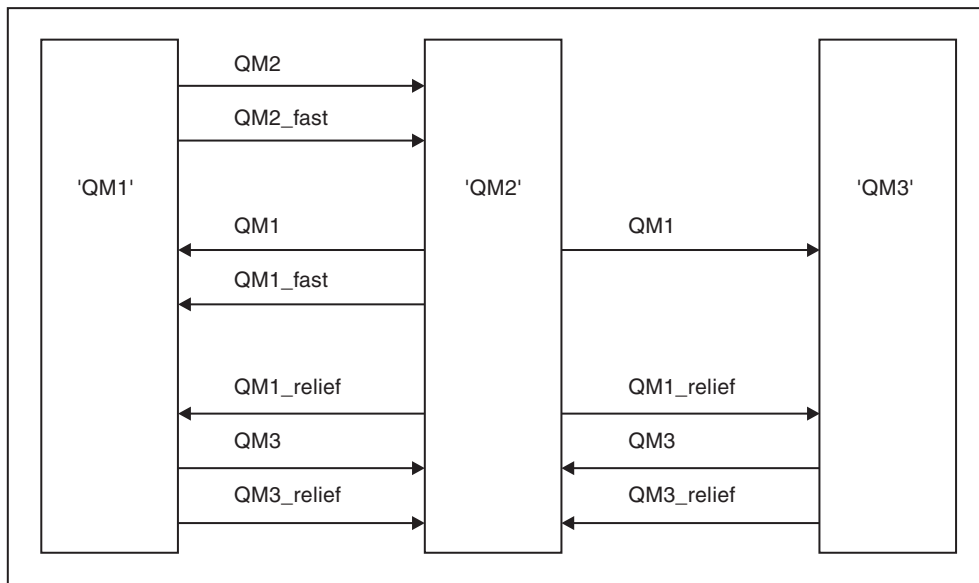


Figure 1. Network diagram showing all channels

Table 1. Example of channel names

Route name	Queue managers hosting channel	Transmission queue name	Suggested channel name
QM1	QM1 & QM2	QM1 (at QM2)	QM1.from.QM2
QM1	QM2 & QM3	QM1 (at QM3)	QM1.from.QM3
QM1_fast	QM1 & QM2	QM1_fast (at QM2)	QM1_fast.from.QM2
QM1_relief	QM1 & QM2	QM1_relief (at QM2)	QM1_relief.from.QM2
QM1_relief	QM2 & QM3	QM1_relief (at QM3)	QM1_relief.from.QM3
QM2	QM1 & QM2	QM2 (at QM1)	QM2.from.QM1
QM2_fast	QM1 & QM2	QM2_fast (at QM1)	QM2_fast.from.QM1
QM3	QM1 & QM2	QM3 (at QM1)	QM3.from.QM1
QM3	QM2 & QM3	QM3 (at QM2)	QM3.from.QM2
QM3_relief	QM1 & QM2	QM3_relief (at QM1)	QM3_relief.from.QM1
QM3_relief	QM2 & QM3	QM3_relief (at QM2)	QM3_relief.from.QM2

Note:

- ▶ **z/OS** On IBM MQ for z/OS, queue-manager names are limited to four characters.

2. Name all the channels in your network uniquely. As shown in Table 1 on page 5, including the source and target queue manager names in the channel name is a good way to do so.

Network planner

Creating a network assumes that there is another, higher level function of *network planner* whose plans are implemented by the other members of the team.

For widely used applications, it is more economical to think in terms of local access sites for the concentration of message traffic, using wide-band links between the local access sites, as shown in Figure 2 on page 7.

In this example there are two main systems and a number of satellite systems. The actual configuration would depend on business considerations. There are two concentrator queue managers located at convenient centers. Each QM-concentrator has message channels to the local queue managers:

- QM-concentrator 1 has message channels to each of the three local queue managers, QM1, QM2, and QM3. The applications using these queue managers can communicate with each other through the QM-concentrators.
- QM-concentrator 2 has message channels to each of the three local queue managers, QM4, QM5, and QM6. The applications using these queue managers can communicate with each other through the QM-concentrators.
- The QM-concentrators have message channels between themselves thus allowing any application at a queue manager to exchange messages with any other application at another queue manager.

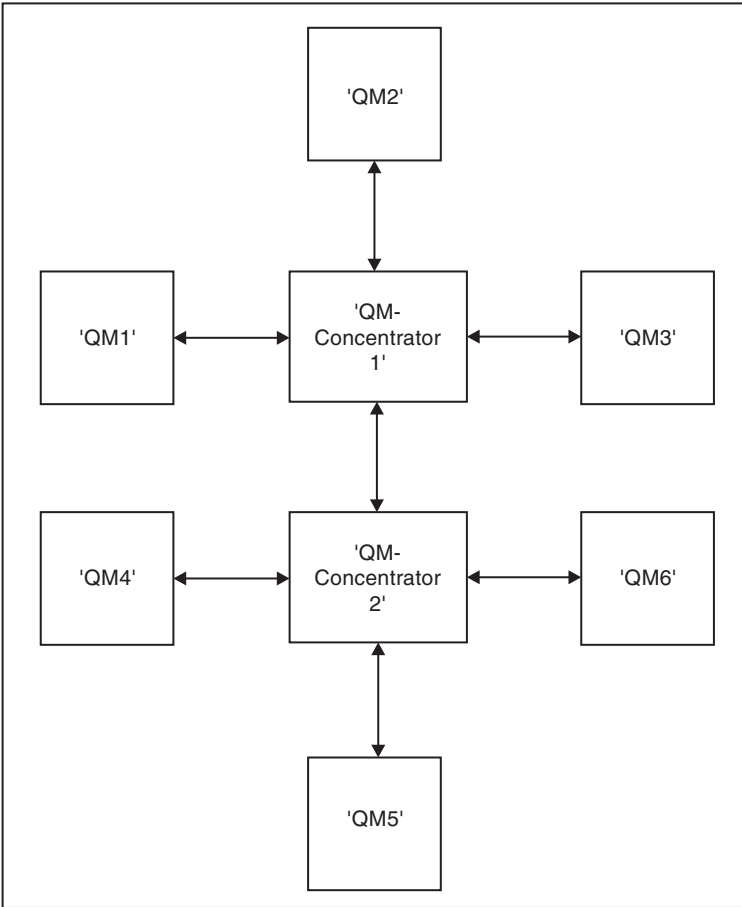


Figure 2. Network diagram showing QM-concentrators

Designing clusters

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. Clusters must be carefully designed, to ensure that they function correctly, and that they achieve the required levels of availability and responsiveness.

Before you begin

For an introduction to clustering concepts, see the following topics:

- Distributed queuing and clusters
- “Comparison of clustering and distributed queuing” on page 14
- Components of a cluster

When you are designing your queue manager cluster you have to make some decisions. You must first decide which queue managers in the cluster are to hold the full repositories of cluster information. Any queue manager you create can work in a cluster. You can choose any number of queue managers for this purpose but the ideal number is two. For information about selecting queue managers to hold the full repositories, see “How to choose cluster queue managers to hold full repositories” on page 16.

See the following topics for more information about designing your cluster:

- “Example clusters” on page 24
- “Organizing a cluster” on page 18
- “Cluster naming conventions” on page 19

- “Queue-sharing groups and clusters” on page 20
- “Overlapping clusters” on page 20

What to do next


See the following topics for more information about configuring and working with clusters:

- Establishing communication in a cluster
- Configuring a queue manager cluster
- Routing messages to and from clusters
- Using clusters for workload management

For more information to help you configure your cluster, see “Clustering tips” on page 21.

Planning how you use multiple cluster transmission queues:

You can explicitly define transmission queues, or have the system generate the transmission queues for you. If you define the transmission queues yourself, you have more control over the queue definitions.

 On z/OS, you also have more control over the page set where the messages are held.

Defining the transmission queues

There are two methods of defining transmission queues:

- Automatically, using the queue manager attribute DEFCLXQ, as follows:
`ALTER QMGR DEFCLXQ(SCTQ | CHANNEL)`
DEFCLXQ(SCTQ) indicates that the default transmission queue for all cluster-sender channels is SYSTEM.CLUSTER.TRANSMIT.QUEUE. This is the default value.
DEFCLXQ(CHANNEL) indicates that by default each cluster-sender channel uses a separate transmission queue named SYSTEM.CLUSTER.TRANSMIT.<channel name>. Each transmission queue is automatically defined by the queue manager. See “Automatically-defined cluster transmission queues” on page 10 for more information.
- Manually, by defining a transmission queue with a value specified for the CLCHNAME attribute. The CLCHNAME attribute indicates which cluster-sender channels should use the transmission queue. See “Planning for manually-defined cluster transmission queues” on page 12 for more information.

What security do I need?

To initiate a switch, either automatically or manually, you need authority to start a channel.

To define the queue used as a transmission queue, you need standard IBM MQ authority to define the queue.

When is a suitable time to implement the change?

When changing the transmission queue used by cluster-sender channels, you need to allocate a time in which to make the update, considering the following points:

- The time required for a channel to switch transmission queue depends on the total number of messages on the old transmission queue, how many messages need to be moved, and the size of the messages.
- Applications can continue to put messages to the transmission queue while the change is happening. This might lead to an increase in the transition time.
- You can change the CLCHNAME parameter of any transmission queue or DEFCLXQ at any time, preferably when the workload is low.

Note that nothing happens immediately.

- Changes occur only when a channel starts or restarts. When a channel starts it checks the current configuration and switches to a new transmission queue if required.
- There are several changes that might alter the association of a cluster-sender channel with a transmission queue:
 - Altering the value of a transmission queue's CLCHNAME attribute, making CLCHNAME less specific or blank.
 - Altering the value of a transmission queue's CLCHNAME attribute, making CLCHNAME more specific.
 - Deleting a queue with CLCHNAME specified.
 - Altering the queue manager attribute DEFCLXQ.

How long will the switch take?

During the transition period, any messages for the channel are moved from one transmission queue to another. The time required for a channel to switch transmission queue depends on the total number of messages on the old transmission queue, and how many messages need to be moved.

For queues containing a few thousand messages, it should take under a second to move the messages. The actual time depends on the number and size of the messages. Your queue manager should be able to move messages at many megabytes each second.

Applications can continue to put messages to the transmission queue while the change is happening. This might lead to an increase in the transition time.

Each affected cluster-sender channel must be restarted for the change to take effect. Therefore, it is best to change the transmission queue configuration when the queue manager is not busy, and few messages are stored on the cluster transmission queues.

The **runswch1** command,  or the SWITCH CHANNEL(*) STATUS command in CSQUTIL on z/OS, can be used to query the status of cluster-sender channels and what pending changes are outstanding to their transmission queue configuration.

How to implement the change

See Implementing the system using multiple cluster transmission queues for details on how you make the change to multiple cluster transmission queues, either automatically or manually.

Undoing the change

See Undoing a change for details on how you back out changes if you encounter problems.

Automatically-defined cluster transmission queues:

You can have the system generate the transmission queues for you.

About this task

If a channel does not have a manually defined cluster transmission queue that is associated with it, and you specify DEFCLXQ(CHANNEL), when the channel starts the queue manager automatically defines a permanent-dynamic queue for the cluster sender channel. This channel has the name SYSTEM.CLUSTER.TRANSMIT.ChannelName, and uses the model queue SYSTEM.CLUSTER.TRANSMIT.MODEL.

z/OS To set up the cluster transmission queues manually, see “Planning for manually-defined cluster transmission queues” on page 12.

Important: **z/OS**

If the queue manager is migrated to IBM MQ Version 8.0, the queue manager does not have the SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE.

Define this queue first, so that the command ALTER QGMR DEFCLXQ(CHANNEL) takes effect.

The following JCL is an example of the code you can use to define the model queue:

```
//CLUSMODL JOB MSGCLASS=H,NOTIFY=&SYSUID
/*JOBPARM SYSAFF=(MVCC)
//MQCMD EXEC PGM=CSQUTIL,REGION=4096K,PARM='CDLK'
//STEPLIB DD DISP=SHR,DSN=SCEN.MQ.V000.COM.BASE.SCSQAUTH
// DD DISP=SHR,DSN=SCEN.MQ.V000.COM.BASE.SCSQANLE
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(CMDINP)
/*
//CMDINP DD *
DEFINE QMODEL( 'SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE' ) +
QSGDISP( QMGR ) +

* COMMON QUEUE ATTRIBUTES
DESCR( 'SYSTEM CLUSTERING TRANSMISSION MODEL QUEUE' ) +
PUT( ENABLED ) +
DEFPRTY( 5 ) +
DEFPSIST( YES ) +

* MODEL QUEUE ATTRIBUTES
DEFTYPE( PERMDYN ) +

* LOCAL QUEUE ATTRIBUTES
GET( ENABLED ) +
SHARE +
DEFSOPT( EXCL ) +
MSGDLVSQ( PRIORITY ) +
RETINTVL( 999999999 ) +
MAXDEPTH( 999999999 ) +
MAXMSGL( 4194304 ) +
NOHARDENBO +
BOTHRESH( 0 ) +
BOQNAME( ' ' ) +
STGCLASS( 'REMOTE' ) +
USAGE( XMITQ ) +
INDXTYPE( CORRELID ) +
CFSTRUCT( ' ' ) +
MONQ( OFF ) ACCTQ( OFF ) +

* EVENT CONTROL ATTRIBUTES
QDPMAEV( ENABLED ) +
QDPHIEV( DISABLED ) +
QDEPTHHI( 80 ) +
QDPLOEV( DISABLED ) +
QDEPTHLO( 40 ) +
QSVCIIEV( NONE ) +
QSVCIINT( 999999999 ) +

* TRIGGER ATTRIBUTES
TRIGGER +
```

```

TRIGTYPE( FIRST ) +
TRIGMPRI( 0 ) +
TRIGDPTH( 1 ) +
TRIGDATA( ' ' ) +
PROCESS( ' ' ) +
INITQ( ' ' )
/*

```

Procedure

1. Use the *DEFCLXQ* queue manager attribute. For more information on this attribute, see ALTER QMGR.

There are two options:

SCTQ This option is the default, and means that you use the single SYSTEM.CLUSTER.TRANSMIT.QUEUE.

CHANNEL

Means that you use multiple cluster transmission queues.


2. To switch to the new association:

- Stop and restart the channel.
- The channel uses the new transmission queue definition.
- Messages are transferred by a transitional switch process from the old queue to the new transmission queue.

Note that any application messages are put to the old definition.

When the number of messages on the old queue reaches zero, new messages are placed directly on the new transmission queue.

3. To monitor when the switching process finishes:

- a. A switch of transmission queue that is initiated by a channel runs in the background and your administrator can monitor the queue manager job log to determine when it has completed.
- b. Monitor messages on the job log to show the progress of the switch.
- c. To make sure that only the channels that you wanted are using this transmission queue, issue the command DIS CLUSQMGR(*) where, for example, the transmission queue property that defines the transmission queue is APPQMGR.CLUSTER1.XMITQ.
- d.  Use the SWITCH CHANNEL (*) STATUS command under CSQUTIL. This option tells you what pending changes are outstanding, and how many messages need to be moved between transmission queues.

Results

You have set up your cluster transmission queue, or queues.

Related tasks:

“Planning for manually-defined cluster transmission queues”

If you define the transmission queues yourself you have more control over the definitions, and the page set on which the messages are held.

Related information:

ALTER QMGR
DISPLAY CLUSQMGR

Planning for manually-defined cluster transmission queues:

If you define the transmission queues yourself you have more control over the definitions, and the page set on which the messages are held.

About this task

Your administrator manually defines a transmission queue and uses a new queue attribute CLCHNAME to define which cluster sender channel, or channels, will use this queue as their transmission queue.

Note that CLCHNAME can include a wild card character at the beginning, or end, to allow a single queue to be used for multiple channels.

To set up cluster transmission queues automatically, see “Automatically-defined cluster transmission queues” on page 10.

Procedure

1. For example, enter the following:

```
DEFINE QLOCAL(APPQMGR.CLUSTER1.XMITQ)
CLCHNAME(CLUSTER1.TO.APPQMGR)
USAGE(XMITQ) STGCLASS(STG1)
INDXTYPE( CORRELID ) SHARE
```

```
DEFINE STGCLASS(STG1) PSID(3)
DEFINE PSID(3) BUFFERPOOL(4)
```

Tip: You need to plan which page set (and buffer pool) you use for your transmission queues. You can have different page sets for different queues, and provide isolation between them, so one page set filling up, does not impact transmission queues in other page sets.

See Working with cluster transmission queues and cluster-sender channels for information on how each channel selects the appropriate queue.

When the channel starts it switches its association to the new transmission queue. In order to make sure no message is lost, the queue manager automatically transfers messages from the old cluster transmission queue to the new transmission queue in order.

2. Use the CSQUTIL SWITCH function to change to the new association. See Switch the transmission queue associated with cluster-sender channels (SWITCH) for further information.
 - a. STOP the channel, or channels, whose transmission queue is to be changed, so that they are in STOPPED status. For example:

```
STOP CHANNEL(CLUSTER1.TO.APPQMGR)
```
 - b. Change the CLCHNAME(XXXX) attribute on the transmission queue.
 - c. Use the SWITCH function to switch the messages or monitor what is happening. Use the command

```
SWITCH CHANNEL(*) MOVEMSGS(YES)
```

to move the messages without starting the channel.

- d. Start the channel, or channels, and check whether the channel is using the correct queues. For example:

```
DIS CHS(CLUSTER1.TO.APPQMGR)
DIS CHS(*) where(XMITQ eq APPQMGR.CLUSTER1.XMITQ)
```

Tip:

- The following process uses the CSQUTIL SWITCH function; for more information, see Switch the transmission queue associated with cluster-sender channels (SWITCH).

You do not have to use this function but using this function gives more options:

- Using SWITCH CHANNEL (*) STATUS provides an easy way to identify the switching status of cluster-sender channels. It allows your administrator to see what channels are currently switching, and those channels having a switch pending that take effect when those channels next start.

Without this capability your administrator needs to use multiple DISPLAY commands, and then process the resulting output to ascertain this information. Your administrator can also confirm that a configuration change has the desired result.

- If CSQUTIL is used to initiate the switch, CSQUTIL continues to monitor the progress of this operation, and only ends when the switch has completed.

This can make it much easier to perform these operations in batch. Also, if CSQUTIL is run to switch multiple channels, CSQUTIL performs these actions sequentially; this can have less impact for your enterprise than multiple switches running in parallel.

Results

You have set up your cluster transmission queue, or queues.

Access control and multiple cluster transmission queues:

Choose between three modes of checking when an application puts messages to remote cluster queues. The modes are checking remotely against the cluster queue, checking locally against SYSTEM.CLUSTER.TRANSMIT.QUEUE, or checking against local profiles for the cluster queue, or cluster queue manager.


IBM MQ gives you the choice of checking locally, or locally and remotely, that a user has permission to put a message to a remote queue. A typical IBM MQ application uses local checking only, and relies on the remote queue manager trusting the access checks made on the local queue manager. If remote checking is not used, the message is put to the target queue with the authority of the remote message channel process. To use remote checking you must set the put authority of the receiving channel to context security.

The local checks are made against the queue that the application opens. In distributed queuing, the application usually opens a remote queue definition, and access checks are made against the remote queue definition. If the message is put with a full routing header, the checks are made against the transmission queue. If an application opens a cluster queue that is not on the local queue manager, there is no local object to check. The access control checks are made against the cluster transmission queue, SYSTEM.CLUSTER.TRANSMIT.QUEUE. Even with multiple cluster transmission queues, from Version 7.5, local access control checks for remote cluster queues are made against SYSTEM.CLUSTER.TRANSMIT.QUEUE.

The choice of local or remote checking is a choice between two extremes. Checking remotely is fine-grained. Every user must have an access control profile on every queue manager in the cluster to put to any cluster queue. Checking locally is coarse-grained. Every user needs only one access control profile for the cluster transmission queue on the queue manager they are connected to. With that profile, they can put a message to any cluster queue on any queue manager in any cluster.

Since Version 7.1, administrators have another way to set up access control for cluster queues. You can create a security profile for a cluster queue on any queue manager in the cluster using the **setmqaut** command. The profile takes effect if you open a remote cluster queue locally, specifying only the queue name. You can also set up a profile for a remote queue manager. If you do so, the queue manager can check the profile of a user that opens a cluster queue by providing a fully qualified name.

The new profiles work only if you change the queue manager stanza, **ClusterQueueAccessControl** to RQMName. The default is Xmi tq. You must create profiles for all the cluster queues existing applications use cluster queues. If you change the stanza to RQMName without creating profiles the applications are likely to fail.

Tip: The changes made to cluster queue accessing checking in Version 7.1 do not apply to remote queuing. Access checks are still made against local definitions. The changes mean that you can follow the same approach to configure access checking on cluster queues and cluster topics.  The changes also align the access checking approach for cluster queues more closely with z/OS. The commands to set up access checking on z/OS are different, but both check access against a profile rather than against the object itself.

Related information:

Clustering: Application isolation using multiple cluster transmission queues
 Setting ClusterQueueAccessControl

Comparison of clustering and distributed queuing:

Compare the components that need to be defined to connect queue managers using distributed queuing and clustering.

If you do not use clusters, your queue managers are independent and communicate using distributed queuing. If one queue manager needs to send messages to another, you must define:

- A transmission queue
- A channel to the remote queue manager

Figure 3 shows the components required for distributed queuing.




Figure 3. Distributed queuing

If you group queue managers in a cluster, queues on any queue manager are available to any other queue manager in the cluster. Any queue manager can send a message to any other queue manager in the same cluster without explicit definitions. You do not provide channel definitions, remote-queue definitions, or transmission queues for each destination. Every queue manager in a cluster has a single transmission queue from which it can transmit messages to any other queue manager in the cluster. Each queue manager in a cluster needs to define only:

- One cluster-receiver channel on which to receive messages

- One cluster-sender channel with which it introduces itself and learns about the cluster

Definitions to set up a cluster versus distributed queuing

Look at Figure 4, which shows four queue managers each with two queues. Consider how many definitions are needed to connect these queue managers using distributed queuing. Compare how many definitions are needed to set up the same network as a cluster.




Figure 4. A network of four queue managers

Definitions to set up a network using distributed queuing

To set up the network shown in Figure 3 on page 14 using distributed queuing, you might have the following definitions:

Table 2. Definitions for distributed queuing

Description	Number per queue manager	Total number
A sender-channel definition for a channel on which to send messages to every other queue manager	3	12
A receiver-channel definition for a channel on which to receive messages from every other queue manager	3	12
A transmission-queue definition for a transmission queue to every other queue manager	3	12
A local-queue definition for each local queue	2	8
A remote-queue definition for each remote queue to which this queue manager wants to put messages	6	24

You might reduce this number of definitions by using generic receiver-channel definitions. The maximum number of definitions could be as many as 17 on each queue manager, which is a total of 68 for this network.

Definitions to set up a network using clusters

To set up the network shown in Figure 3 on page 14 using clusters you need the following definitions:

Table 3. Definitions for clustering

Description	Number per queue manager	Total number
A cluster-sender channel definition for a channel on which to send messages to a repository queue manager	1	4
A cluster-receiver channel definition for a channel on which to receive messages from other queue managers in the cluster	1	4
A local-queue definition for each local queue	2	8

To set up this cluster of queue managers (with two full repositories), you need four definitions on each queue manager, a total of sixteen definitions altogether. You also need to alter the queue-manager definitions for two of the queue managers, to make them full repository queue managers for the cluster.

Only one CLUSSDR and one CLUSRCVR channel definition is required. When the cluster is defined, you can add or remove queue managers (other than the repository queue managers) without any disruption to the other queue managers.

Using a cluster reduces the number of definitions required to set up a network containing many queue managers.

With fewer definitions to make there is less risk of error:

- Object names always match, for example the channel name in a sender-receiver pair.
- The transmission queue name specified in a channel definition always matches the correct transmission queue definition or the transmission queue name specified in a remote queue definition.
- A QREMOTE definition always points to the correct queue at the remote queue manager.

Once a cluster is set up, you can move cluster queues from one queue manager to another within the cluster without having to do any system management work on any other queue manager. There is no chance of forgetting to delete or modify channel, remote-queue, or transmission-queue definitions. You can add new queue managers to a cluster without any disruption to the existing network.

How to choose cluster queue managers to hold full repositories:

In each cluster you must choose at least one, and preferably two queue managers to hold full repositories. Two full repositories are sufficient for all but the most exceptional circumstances. If possible, choose queue managers that are hosted on robust and permanently-connected platforms, that do not have coinciding outages, and that are in a central position geographically. Also consider dedicating systems as full repository hosts, and not using these systems for any other tasks.

Full repositories are queue managers that maintain a complete picture of the state of the cluster. To share this information, each full repository is connected by CLUSSDR channels (and their corresponding CLUSRCVR definitions) to every other full repository in the cluster. You must manually define these channels.




Figure 5. Two connected full repositories.

Every other queue manager in the cluster maintains a picture of what it currently knows about the state of the cluster in a *partial repository*. These queue managers publish information about themselves, and

request information about other queue managers, using any two available full repositories. If a chosen full repository is not available, another is used. When the chosen full repository becomes available again, it collects the latest new and changed information from the others so that they keep in step. If all the full repositories go out of service, the other queue managers use the information they have in their partial repositories. However, they are limited to using the information that they have; new information and requests for updates cannot be processed. When the full repositories reconnect to the network, messages are exchanged to bring all repositories (both full and partial) up to date.

When planning the allocation of full repositories, include the following considerations:

- The queue managers chosen to hold full repositories need to be reliable and managed. Choose queue managers that are hosted on a robust and permanently-connected platform.
- Consider the planned outages for the systems hosting your full repositories, and ensure that they do not have coinciding outages.
- Consider network performance: Choose queue managers that are in a central position geographically, or that share the same system as other queue managers in the cluster.
- Consider whether a queue manager is a member of more than one cluster. It can be administratively convenient to use the same queue manager to host the full repositories for several clusters, provided this benefit is balanced against how busy you expect the queue manager to be.
- Consider dedicating some systems to contain only full repositories, and not using these systems for any other tasks. This ensures that these systems only require maintenance for queue manager configuration, and are not removed from service for the maintenance of other business applications. It also ensures that the task of maintaining the repository does not compete with applications for system resources. This can be particularly beneficial in large clusters (say, clusters of more than a thousand queue managers), where full repositories have a much higher workload in maintaining the cluster state.

Having more than two full repositories is possible, but rarely recommended. Although object definitions (that is, queues, topics and channels) flow to all available full repositories, requests only flow from a partial repository to a maximum of two full repositories. This means that, when more than two full repositories are defined, and any two full repositories become unavailable, some partial repositories might not receive updates they would expect. See *WMQ Clusters: Why only two Full Repositories?*

One situation in which you might find it useful to define more than two full repositories is when migrating existing full repositories to new hardware or new queue managers. In this case, you should introduce the replacement full repositories, and confirm that they have become fully populated, before you remove the previous full repositories. Whenever you add a full repository, remember that you must directly connect it to every other full repository with CLUSSDR channels.




Figure 6. More than two connected full repositories

Related information:

- 🔗 [WMQ Clusters: Why only two Full Repositories?](#)
- 🔗 [How big can a WMQ Cluster be?](#)

Organizing a cluster:

Select which queue managers to link to which full repository. Consider the performance effect, the queue manager version, and whether multiple CLUSSDR channels are desirable.


Having selected the queue managers to hold full repositories, you need to decide which queue managers to link to which full repository. The CLUSSDR channel definition links a queue manager to a full repository from which it finds out about the other full repositories in the cluster. From then on, the queue manager sends messages to any two full repositories. It always tries to use the one to which it has a CLUSSDR channel definition first. You can choose to link a queue manager to either full repository. In choosing, consider the topology of your configuration, and the physical or geographical location of the queue managers.

▶ **z/OS** Queue managers on IBM MQ z/OS Version 6.0 or later support PCF commands. The MQ Explorer can use them to query information about the cluster (in just the same way as with distributed queue managers).

▶ **z/OS** z/OS queue managers up to IBM MQ v5.3.1 do not support PCF commands. The MQ Explorer cannot use them to query information about the cluster. For this reason, the MQ Explorer does not allow you to select IBM MQ v5.3.1 z/OS queue managers at the source of the information about a cluster.

Because all cluster information is sent to two full repositories, there might be situations in which you want to make a second CLUSSDR channel definition. You might define a second CLUSSDR channel in a cluster that has many full repositories spread over a wide area. You can then control which two full repositories your information is sent to.

Cluster naming conventions:

Consider naming queue managers in the same cluster using a naming convention that identifies the cluster to which the queue manager belongs. Use a similar naming convention for channel names and extend it to describe the channel characteristics.  Do not use a generic connection in defining cluster-receiver channels on z/OS.

This information contains the old guidance on naming conventions, and the current guidance. As the IBM MQ technology improves, and as customers use technology in new or different ways, new recommendations and information must be provided for these scenarios.

Cluster naming conventions: Old best practices


When setting up a new cluster, consider a naming convention for the queue managers. Every queue manager must have a different name. If you give queue managers in a cluster a set of similar names, it might help you to remember which queue managers are grouped where.

When defining channels, remember that all cluster-sender channels have the same name as their corresponding cluster-receiver channel. Channel names are limited to a maximum of 20 characters.

Every cluster-receiver channel must also have a unique name. One possibility is to use the *queue-manager* name preceded by the *cluster-name*. For example, if the *cluster-name* is CLUSTER1 and the *queue-managers* are QM1, QM2, then cluster-receiver channels are CLUSTER1.QM1, CLUSTER1.QM2.

You might extend this convention if channels have different priorities or use different protocols; for example, CLUSTER1.QM1.S1, CLUSTER1.QM1.N3, and CLUSTER1.QM1.T4. In this example, S1 might be the first SNA channel, N3 might be the NetBIOS channel with a network priority of three.

A final qualifier might describe the class of service the channel provides.

 In IBM MQ for z/OS, you can define VTAM generic resources or *Dynamic Domain Name Server* (DDNS) generic names. You can define connection names using generic names. However, when you create a cluster-receiver definition, do not use a generic connection name.

The problem with using generic connection names for cluster-receiver definitions is as follows. If you define a CLUSRCVR with a generic CONNAME there is no guarantee that your CLUSSDR channels point to the queue managers you intend. Your initial CLUSSDR might end up pointing to any queue manager in the queue-sharing group, not necessarily one that hosts a full repository. If a channel starts trying a connection again, it might reconnect to a different queue manager with the same generic name disrupting the flow of messages.

Cluster naming conventions: Current best practices

A good naming convention can help to reduce confusion over ownership and scope of clusters. A clear naming convention throughout the cluster topology causes a lot less confusion if clusters get merged at a later time. This situation is also improved if everyone involved is clear about who owns which queue managers and which clusters. Probably the most important point for cluster naming conventions is to put the queue manager name in the channel name, see the following example:

```
CLUSNAME.QMGRNAME
```

This convention might not be obvious to experienced IBM MQ users that are unfamiliar with clusters. This oversight is because the XXX.T0.YYY format is such a common method. For example, CLUSTER.T0.XX or CLUSTER.X are commonly used formats that are not recommended for clustering, as they can quickly

reach the 20 character limit. The commonly used `CLUSTER.TO.XX` format becomes confusing if another channel is added later (for example when joining another cluster).

Other objects also benefit from sensible rules, such as: `LOB.PROJECT.QNAME` or `LOB.CLUSTER.ALIAS.NAME`.

Queue-sharing groups and clusters:

Shared queues can be cluster queues and queue managers in a queue-sharing group can also be cluster queue managers.

▶ **z/OS** On IBM MQ for z/OS you can group queue managers into queue-sharing groups. A queue manager in a queue-sharing group can define a local queue that is to be shared by up to 32 queue managers.

Shared queues can also be cluster queues. Furthermore, the queue managers in a queue-sharing group can also be in one or more clusters.

▶ **z/OS** In IBM MQ for z/OS, you can define VTAM generic resources or *Dynamic Domain Name Server* (DDNS) generic names. You can define connection names using generic names. However, when you create a cluster-receiver definition, do not use a generic connection name.

▶ **z/OS** The problem with using generic connection names for cluster-receiver definitions is as follows. If you define a `CLUSRCVR` with a generic `CONNNAME` there is no guarantee that your `CLUSSDR` channels point to the queue managers you intend. Your initial `CLUSSDR` might end up pointing to any queue manager in the queue-sharing group, not necessarily one that hosts a full repository. If a channel starts trying a connection again, it might reconnect to a different queue manager with the same generic name disrupting the flow of messages.

Overlapping clusters:

Overlapping clusters provide additional administrative capabilities. Use namelists to reduce the number of commands needed to administer overlapping clusters.

You can create clusters that overlap. There are a number of reasons you might define overlapping clusters; for example:

- To allow different organizations to have their own administration.
- To allow independent applications to be administered separately.
- To create classes of service.

In Figure 7 on page 21, the queue manager `STF2` is a member of both the clusters. When a queue manager is a member of more than one cluster, you can take advantage of namelists to reduce the number of definitions you need. Namelists contain a list of names, for example, cluster names. You can create a namelist naming the clusters. Specify the namelist on the `ALTER QMGR` command for `STF2` to make it a full repository queue manager for both clusters.

If you have more than one cluster in your network, you must give them different names. If two clusters with the same name are ever merged, it is not possible to separate them again. It is also a good idea to give the clusters and channels different names. They are more easily distinguished when you look at the output from the `DISPLAY` commands. Queue manager names must be unique within a cluster for it to work correctly.

Defining classes of service

Imagine a university that has a queue manager for each member of staff and each student. Messages between members of staff are to travel on channels with a high priority and high bandwidth. Messages

between students are to travel on cheaper, slower channels. You can set up this network using traditional distributed queuing techniques. IBM MQ selects which channels to use by looking at the destination queue name and queue manager name.

To clearly differentiate between the staff and students, you could group their queue managers into two clusters as shown in Figure 7. IBM MQ moves messages to the meetings queue in the staff cluster only over channels that are defined in that cluster. Messages for the gossip queue in the students cluster go over channels defined in that cluster and receive the appropriate class of service.




Figure 7. Classes of service

Clustering tips:

You might need to make some changes to your systems or applications before using clustering. There are both similarities and differences from the behavior of distributed queuing.

- **z/OS** The IBM MQ Explorer cannot directly administer IBM MQ for z/OS queue managers at versions earlier than Version 6.0.
- You must add manual configuration definitions to queue managers outside a cluster for them to access cluster queues.
- If you merge two clusters with the same name, you cannot separate them again. Therefore it is advisable to give all clusters a unique name.
- If a message arrives at a queue manager but there is no queue there to receive it, the message is put on the dead-letter queue. If there is no dead-letter queue, the channel fails and tries again. The use of the dead-letter queue is the same as with distributed queuing.
- The integrity of persistent messages is maintained. Messages are not duplicated or lost as a result of using clusters.
- Using clusters reduces system administration. Clusters make it easy to connect larger networks with many more queue managers than you would be able to contemplate using distributed queuing. There is a risk that you might consume excessive network resources if you attempt to enable communication between every queue manager in a cluster.
- If you use the IBM MQ Explorer, which presents the queue managers in a tree structure, the view for large clusters might be cumbersome.
- **z/OS** IBM MQ Explorer can administer a cluster with repository queue managers on IBM MQ for z/OS Version 6 or later. You need not nominate an additional repository on a separate system. For earlier versions of IBM MQ on z/OS, the IBM MQ Explorer cannot administer a cluster with repository queue managers. You must nominate an additional repository on a system that the IBM MQ Explorer can administer.

- The purpose of distribution lists is to use a single MQPUT command to send the same message to multiple destinations. Distribution lists are supported on IBM MQ for AIX®, IBM i, HP-UX, Solaris, Linux, and Windows. You can use distribution lists with queue manager clusters. In a cluster, all the messages are expanded at MQPUT time. The advantage, in terms of network traffic, is not so great as in a non-clustering environment. The advantage of distribution lists is that the numerous channels and transmission queues do not need to be defined manually.
- If you are going to use clusters to balance your workload examine your applications. See whether they require messages to be processed by a particular queue manager or in a particular sequence. Such applications are said to have message affinities. You might need to modify your applications before you can use them in complex clusters.
- You might choose to use the MQ00_BIND_ON_OPEN option on an MQOPEN to force messages to be sent to a specific destination. If the destination queue manager is not available the messages are not delivered until the queue manager becomes available again. Messages are not routed to another queue manager because of the risk of duplication.
- If a queue manager is to host a cluster repository, you need to know its host name or IP address. You have to specify this information in the CONNAME parameter when you make the CLUSSDR definition on other queue managers joining the cluster. If you use DHCP, the IP address is subject to change because DHCP can allocate a new IP address each time you restart a system. Therefore, you must not specify the IP address in the CLUSSDR definitions. Even if all the CLUSSDR definitions specify the host name rather than the IP address, the definitions would still not be reliable. DHCP does not necessarily update the DNS directory entry for the host with the new address. If you must nominate queue managers as full repositories on systems that use DHCP, install software that guarantees to keep your DNS directory up to date.
- Do not use generic names, for example VTAM generic resources or Dynamic Domain Name Server (DDNS) generic names as the connection names for your channels. If you do, your channels might connect to a different queue manager than expected.
- You can only get a message from a local cluster queue, but you can put a message to any queue in a cluster. If you open a queue to use the MQGET command, the queue manager opens the local queue.
- You do not need to alter any of your applications if you set up a simple IBM MQ cluster. The application can name the target queue on the MQOPEN call and does not need to know about the location of the queue manager. If you set up a cluster for workload management you must review your applications and modify them as necessary.
- You can view current monitoring and status data for a channel or queue using the DISPLAY CHSTATUS and the DISPLAY QSTATUS **runmqsc** commands. The monitoring information can be used to help gauge the performance and health of the system. Monitoring is controlled by queue manager, queue, and channel attributes. Monitoring of auto-defined cluster-sender channels is possible with the MONACLS queue manager attribute.

Related concepts:

“Comparison of clustering and distributed queuing” on page 14

Compare the components that need to be defined to connect queue managers using distributed queuing and clustering.

Related information:

Clusters

Configuring a queue manager cluster

Components of a cluster

Setting up a new cluster

How long do the queue manager repositories retain information?:

Queue manager repositories retain information for 30 days. An automatic process efficiently refreshes information that is being used.

When a queue manager sends out some information about itself, the full and partial repository queue managers store the information for 30 days. Information is sent out, for example, when a queue manager advertises the creation of a new queue. To prevent this information from expiring, queue managers automatically resend all information about themselves after 27 days. If a partial repository sends a new request for information part way through the 30 day lifetime, the expiry time remains the original 30 days.

When information expires, it is not immediately removed from the repository. Instead it is held for a grace period of 60 days. If no update is received within the grace period, the information is removed. The grace period allows for the fact that a queue manager might have been temporarily out of service at the expiry date. If a queue manager becomes disconnected from a cluster for more than 90 days, it stops being part of the cluster. However, if it reconnects to the network it becomes part of the cluster again. Full repositories do not use information that has expired to satisfy new requests from other queue managers.

Similarly, when a queue manager sends a request for up-to-date information from a full repository, the request lasts for 30 days. After 27 days IBM MQ checks the request. If it has been referenced during the 27 days, it is refreshed automatically. If not, it is left to expire and is refreshed by the queue manager if it is needed again. Expiring requests prevents a buildup of requests for information from dormant queue managers.

Note: For large clusters, it can be disruptive if many queue managers automatically resend all information about themselves at the same time. See Refreshing in a large cluster can affect performance and availability of the cluster.

Related information:

Clustering: Using REFRESH CLUSTER best practices

Example clusters:

The first example shows the smallest possible cluster of two queue managers. The second and third examples show two versions of a three queue manager cluster.

The smallest possible cluster contains only two queue managers. In this case both queue managers contain full repositories. You need only a few definitions to set up the cluster, and yet there is a high degree of autonomy at each queue manager.

DEMOCLSTR

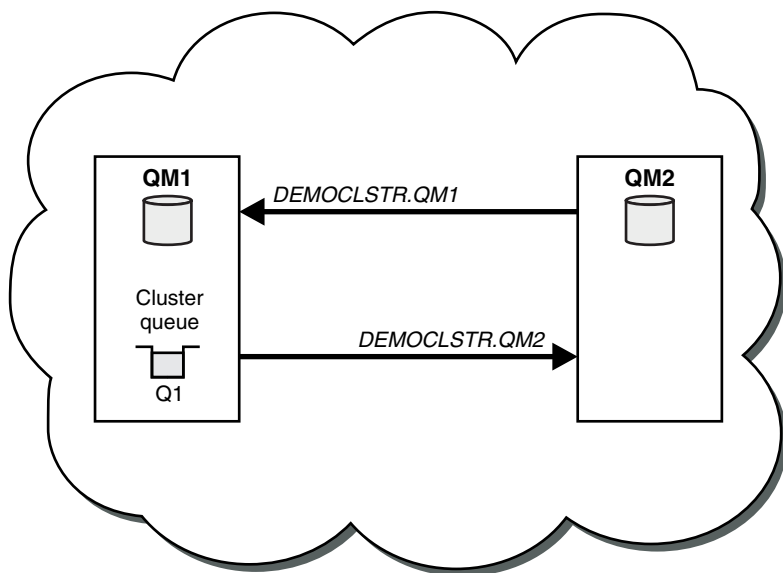


Figure 8. A small cluster of two queue managers

- The queue managers can have long names such as LONDON and NEWYORK. z/OS On IBM MQ for z/OS, queue-manager names are limited to four characters.
- Each queue manager is typically configured on a separate machine. However, you can have multiple queue managers on the same machine.

For instructions on setting up a similar example cluster, see [Setting up a new cluster](#).

Figure 9 on page 25 shows the components of a cluster called CLSTR1.

- In this cluster, there are three queue managers, QM1, QM2, and QM3.
- QM1 and QM2 host repositories of information about all the queue managers and cluster-related objects in the cluster. They are referred to as *full repository queue managers*. The repositories are represented in the diagram by the shaded cylinders.
- QM2 and QM3 host some queues that are accessible to any other queue manager in the cluster. Queues that are accessible to any other queue manager in the cluster are called *cluster queues*. The cluster queues are represented in the diagram by the shaded queues. Cluster queues are accessible from anywhere in the cluster. IBM MQ clustering code ensures that remote queue definitions for cluster queues are created on any queue manager that refers to them.

As with distributed queuing, an application uses the MQPUT call to put a message on a cluster queue at any queue manager in the cluster. An application uses the MQGET call to retrieve messages from a cluster queue only on the queue manager where the queue resides.

- Each queue manager has a manually created definition for the receiving end of a channel called *cluster-name.queue-manager* on which it can receive messages. On the receiving queue manager, *cluster-name.queue-manager* is a cluster-receiver channel. A cluster-receiver channel is like a receiver channel used in distributed queuing; it receives messages for the queue manager. In addition, it also receives information about the cluster.
-




Figure 9. A cluster of queue managers

- In Figure 10 on page 26 each queue manager also has a definition for the sending end of a channel. It connects to the cluster-receiver channel of one of the full repository queue managers. On the sending queue manager, *cluster-name.queue-manager* is a cluster-sender channel. QM1 and QM3 have cluster-sender channels connecting to CLSTR1.QM2, see dotted line "2". QM2 has a cluster-sender channel connecting to CLSTR1.QM1, see dotted line "3". A cluster-sender channel is like a sender-channel used in distributed queuing; it sends messages to the receiving queue manager. In addition, it also sends information about the cluster. Once both the cluster-receiver end and the cluster-sender end of a channel are defined, the channel starts automatically.

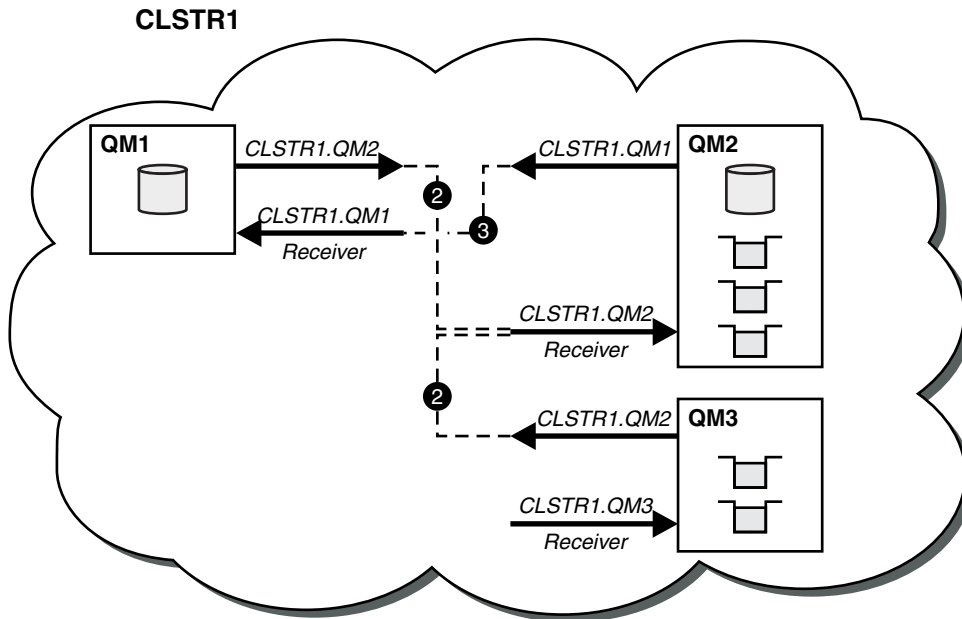


Figure 10. A cluster of queue managers with sender channels

Defining a cluster-sender channel on the local queue manager introduces that queue manager to one of the full repository queue managers. The full repository queue manager updates the information in its full repository accordingly. Then it automatically creates a cluster-sender channel back to the original queue manager, and sends that queue manager information about the cluster. Thus a queue manager learns about a cluster and a cluster learns about a queue manager.

Look again at Figure 9 on page 25. Suppose that an application connected to queue manager QM3 wants to send some messages to the queues at QM2. The first time that QM3 must access those queues, it discovers them by consulting a full repository. The full repository in this case is QM2, which is accessed using the sender channel CLSTR1.QM2. With the information from the repository, it can automatically create remote definitions for those queues. If the queues are on QM1, this mechanism still works, because QM2 is a full repository. A full repository has a complete record of all the objects in the cluster. In this latter case, QM3 would also automatically create a cluster-sender channel corresponding to the cluster-receiver channel on QM1, allowing direct communication between the two.

Figure 11 on page 27 shows the same cluster, with the two cluster-sender channels that were created automatically. The cluster-sender channels are represented by the two dashed lines that join with the cluster-receiver channel CLSTR1.QM3. It also shows the cluster transmission queue, SYSTEM.CLUSTER.TRANSMIT.QUEUE, which QM1 uses to send its messages. All queue managers in the cluster have a cluster transmission queue, from which they can send messages to any other queue manager in the same cluster.

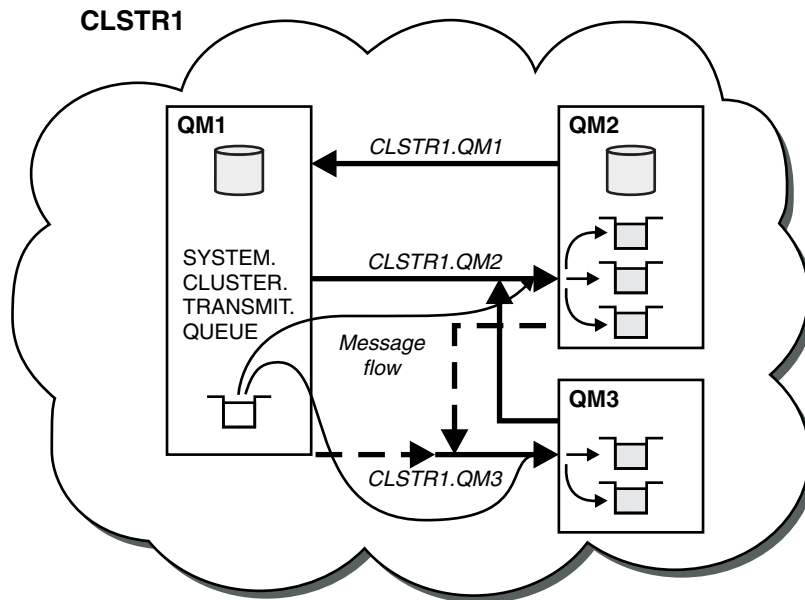


Figure 11. A cluster of queue managers, showing auto-defined channels

Note: Other diagrams show only the receiving ends of channels for which you make manual definitions. The sending ends are omitted because they are mostly defined automatically when needed. The auto-definition of most cluster-sender channels is crucial to the function and efficiency of clusters.

Related concepts:

“Comparison of clustering and distributed queuing” on page 14

Compare the components that need to be defined to connect queue managers using distributed queuing and clustering.

Related information:

Configuring a queue manager cluster

Components of a cluster

Setting up a new cluster

Clustering: Best practices:

Clusters provide a mechanism for interconnecting queue managers. The best practices described in this section are based on testing and feedback from customers.

A successful cluster setup is dependent on good planning and a thorough understanding of IBM MQ fundamentals, such as good application management and network design. Ensure that you are familiar with the information in the related topics before continuing.

Related tasks:

“Designing clusters” on page 7

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. Clusters must be carefully designed, to ensure that they function correctly, and that they achieve the required levels of availability and responsiveness.

Related information:

Distributed queuing and clusters

Clusters

Monitoring clusters

Clustering: Special considerations for overlapping clusters:

This topic provides guidance for planning and administering IBM MQ clusters. This information is a guide based on testing and feedback from customers.

- “Cluster ownership”
- “Overlapping clusters: Gateways”
- “Cluster naming conventions” on page 29

Cluster ownership

Familiarize yourself with overlapping clusters before reading the following information. See “Overlapping clusters” on page 20 and Configuring message paths between clusters for the necessary information.

When configuring and managing a system that consists of overlapping clusters, it is best to adhere to the following:

- Although IBM MQ clusters are 'loosely coupled' as previously described, it is useful to consider a cluster as a single unit of administration. This concept is used because the interaction between definitions on individual queue managers is critical to the smooth functioning of the cluster. For example: When using workload balanced cluster queues it is important that a single administrator or team understand the full set of possible destinations for messages, which depends on definitions spread throughout the cluster. More trivially, cluster sender/receiver channel pairs must be compatible throughout.
- Considering this previous concept; where multiple clusters meet (which are to be administered by separate teams / individuals), it is important to have clear policies in place controlling administration of the gateway queue managers.
- It is useful to treat overlapping clusters as a single namespace: Channel names and queue manager names must be unique throughout a single cluster. Administration is much easier when unique throughout the entire topology. It is best to follow a suitable naming convention, possible conventions are described in “Cluster naming conventions” on page 29.
- Sometimes administrative and system management cooperation is essential/unavoidable: For example, cooperation between organizations that own different clusters that need to overlap. A clear understanding of who owns what, and enforceable rules/conventions helps clustering run smoothly when overlapping clusters.

Overlapping clusters: Gateways

In general, a single cluster is easier to administer than multiple clusters. Therefore creating large numbers of small clusters (one for every application for example) is something to be avoided generally.

However, to provide classes of service, you can implement overlapping clusters. For example:

- Concentric clusters where the smaller one is for Publish/Subscribe. See How to size systems: for more information.

- Some queue managers are to be administered by different teams (see “Cluster ownership” on page 28).
- If it makes sense from an organizational or geographical point of view.
- Equivalent clusters to work with name resolution, as when implementing SSL (or TLS) in an existing cluster.

There is no security benefit from overlapping clusters; allowing clusters administered by two different teams to overlap, effectively joins the teams as well as the topology. Any:

- Name advertised in such a cluster is accessible to the other cluster.
- Name advertised in one cluster can be advertised in the other to draw off eligible messages.
- Non-advertised object on a queue manager adjacent to the gateway can be resolved from any clusters of which the gateway is a member.

The namespace is the union of both clusters and must be treated as a single namespace. Therefore, ownership of an overlapping cluster is shared amongst all the administrators of both clusters.

When a system contains multiple clusters, there might be a requirement to route messages from queue managers in one cluster to queues on queue managers in another cluster. In this situation, the multiple clusters must be interconnected in some way: A good pattern to follow is the use of gateway queue managers between clusters. This arrangement avoids building up a difficult-to-manage mesh of point-to-point channels, and provides a good place to manage such issues as security policies. There are two distinct ways of achieving this arrangement:

1. Place one (or more) queue managers in both clusters using a second cluster receiver definition. This arrangement involves fewer administrative definitions but, as previously stated, means that ownership of an overlapping cluster is shared amongst all the administrators of both clusters.
2. Pair a queue manager in cluster one with a queue manager in cluster two using traditional point-to-point channels.

In either of these cases, various tools can be used to route traffic appropriately. In particular, queue or queue manager aliases can be used to route into the other cluster, and a queue manager alias with blank **RQMNAME** property re-drives workload balancing where it is wanted.

Cluster naming conventions

This information contains the previous guidance on naming conventions, and the current guidance. As the IBM MQ technology improves, and as customers use technology in new or different ways, new recommendations and information must be provided for these scenarios.

Cluster naming conventions: Previous guidance

When setting up a new cluster, consider a naming convention for the queue managers. Every queue manager must have a different name, but it might help you to remember which queue managers are grouped where if you give them a set of similar names.

Every cluster-receiver channel must also have a unique name.

If you have more than one channel to the same queue manager, each with different priorities or using different protocols, you might extend the names to include the different protocols; for example QM1.S1, QM1.N3, and QM1.T4. In this example, S1 might be the first SNA channel, N3 might be the NetBIOS channel with a network priority of 3.

The final qualifier might describe the class of service the channel provides. For more information, see Defining classes of service.

Remember that all cluster-sender channels have the same name as their corresponding cluster-receiver channel.

Do not use generic connection names on your cluster-receiver definitions. In IBM MQ for z/OS, you can define VTAM generic resources or *Dynamic Domain Name Server* (DDNS) generic names, but do not do this if you are using clusters. If you define a CLUSRCVR with a generic **CONNNAME**, there is no guarantee that your CLUSSDR channels point to the queue managers that you intend. Your initial CLUSSDR might end up pointing to any queue manager in the queue-sharing group, not necessarily one that hosts a full repository. Furthermore, if a channel goes to retry status, it might reconnect to a different queue manager with the same generic name and the flow of your messages is disrupted.

Cluster naming conventions: Current guidance

The previous guidance in the section, “Cluster naming conventions: Previous guidance” on page 29, is still valid. However the following guidance is intended as an update when designing new clusters. This updated suggestion ensures uniqueness of channels across multiple clusters, allowing multiple clusters to be successfully overlapped. Because queue managers and clusters can have names of up to 48 characters, and a channel name is limited to 20 characters, care must be taken when naming objects from the beginning to avoid having to change the naming convention midway through a project.

When setting up a new cluster, consider a naming convention for the queue managers. Every queue manager must have a different name. If you give queue managers in a cluster a set of similar names, it might help you to remember which queue managers are grouped where.

When defining channels, remember that all automatically created cluster-sender channels on any queue manager in the cluster have the same name as their corresponding cluster-receiver channel configured on the receiving queue manager in the cluster, and must therefore be unique and make sense across the cluster to the administrators of that cluster. Channel names are limited to a maximum of 20 characters.

One possibility is to use the queue-manager name preceded by the cluster-name. For example, if the cluster-name is CLUSTER1 and the queue-managers are QM1, QM2, then cluster-receiver channels are CLUSTER1.QM1, CLUSTER1.QM2.

You might extend this convention if channels have different priorities or use different protocols; for example, CLUSTER1.QM1.S1, CLUSTER1.QM1.N3, and CLUSTER1.QM1.T4. In this example, S1 might be the first SNA channel, N3 might be the NetBIOS channel with a network priority of three.

A final qualifier might describe the class of service the channel provides.

► **z/OS** In IBM MQ for z/OS, you can define VTAM generic resources or *Dynamic Domain Name Server* (DDNS) generic names. You can define connection names using generic names. However, when you create a cluster-receiver definition, do not use a generic connection name.

► **z/OS** The problem with using generic connection names for cluster-receiver definitions is as follows. If you define a CLUSRCVR with a generic CONNAME there is no guarantee that your CLUSSDR channels point to the queue managers you intend. Your initial CLUSSDR might end up pointing to any queue manager in the queue-sharing group, not necessarily one that hosts a full repository. If a channel starts trying a connection again, it might reconnect to a different queue manager with the same generic name disrupting the flow of messages.

Clustering: Topology design considerations:

This topic provides guidance for planning and administering IBM MQ clusters. This information is a guide based on testing and feedback from customers.

By thinking about where user applications and internal administrative processes are going to be located in advance, many problems can either be avoided, or minimized at a later date. This topic contains information about design decisions that can improve performance, and simplify maintenance tasks as the cluster scales.

- “Performance of the clustering infrastructure”
- “Full repositories” on page 32
- “Should applications use queues on full repositories?” on page 33
- “Managing channel definitions” on page 33
- “Workload balancing over multiple channels” on page 33

Performance of the clustering infrastructure

When an application tries to open a queue on a queue manager in a cluster, the queue manager registers its interest with the full repositories for that queue so that it can learn where the queue exists in the cluster. Any updates to the queue location or configuration are automatically sent by the full repositories to the interested queue manager. This registering of interest is internally known as a subscription (these subscriptions are not the same as IBM MQ subscriptions used for publish/subscribe messaging in IBM MQ)

All information about a cluster goes through every full repository. Full repositories are therefore always being used in a cluster for administrative message traffic. The high usage of system resources when managing these subscriptions, and the transmission of them and the resulting configuration messages, can cause a considerable load on the clustering infrastructure. There are a number of things to consider when ensuring that this load is understood and minimized wherever possible:

- The more individual queue managers using a cluster queue, the more subscriptions are in the system, and thus the bigger the administrative overhead when changes occur and interested subscribers need to be notified, especially on the full repository queue managers. One way to minimize unnecessary traffic and full repository load is by connecting similar applications (that is, those applications that work with the same queues) to a smaller number of queue managers.
- In addition to the number of subscriptions in the system affecting the performance the rate of change in the configuration of clustered objects can affect performance, for example the frequent changing of a clustered queue configuration.
- When a queue manager is a member of multiple clusters (that is, it is part of an overlapping cluster system) any interest made in a queue results in a subscription for each cluster it is a member of, even if the same queue managers are the full repositories for more than one of the clusters. This arrangement increases the load on the system, and is one reason to consider whether multiple overlapping clusters are necessary, rather than a single cluster.
- Application message traffic (that is, the messages being sent by IBM MQ applications to the cluster queues) does not go via the full repositories to reach the destination queue managers. This message traffic is sent directly between the queue manager where the message enters the cluster, and the queue manager where the cluster queue exists. It is not therefore necessary to accommodate high rates of application message traffic with respect to the full repository queue managers, unless the full repository queue managers happen to be either of those two queue managers mentioned. For that reason, it is recommended that full repository queue managers are not used for application message traffic in clusters where the clustering infrastructure load is significant.

Full repositories

A repository is a collection of information about the queue managers that are members of a cluster. A queue manager that hosts a complete set of information about every queue manager in the cluster has a full repository. For more information about full repositories and partial repositories, see Cluster repository.

Full repositories must be held on servers that are reliable and as highly available as possible and single points of failure must be avoided. The cluster design must always have two full repositories. If there is a failure of a full repository, the cluster can still operate.

Details of any updates to cluster resources made by a queue manager in a cluster; for example, clustered queues, are sent from that queue manager to two full repositories at most in that cluster (or to one if there is only one full repository queue manager in the cluster). Those full repositories hold the information and propagate it to any queue managers in the cluster that show an interest in it (that is, they subscribe to it). To ensure that each member of the cluster has an up-to-date view of the cluster resources there, each queue manager must be able to communicate with at least one full repository queue manager at any one time.

If, for any reason a queue manager cannot communicate with any full repositories, it can continue to function in the cluster based on its already cached level of information for a period time, but no new updates or access to previously unused cluster resources are available.

For this reason, you must aim to keep the two full repositories available at all times. However, this arrangement does not mean that extreme measures must be taken because the cluster functions adequately for a short while without a full repository.

There is another reason that a cluster must have two full repository queue managers, other than the availability of cluster information: This reason is to ensure that the cluster information held in the full repository cache exists in two places for recovery purposes. If there is only one full repository, and it loses its information about the cluster, then manual intervention on all queue managers within the cluster is required in order to get the cluster working again. If there are two full repositories however, then because information is always published to and subscribed for from two full repositories, the failed full repository can be recovered with the minimum of effort.

- It is possible to perform maintenance on full repository queue managers in a two full repository cluster design without impacting users of that cluster: The cluster continues to function with only one repository, so where possible bring the repositories down, apply the maintenance, and back up again one at a time. Even if there is an outage on the second full repository, running applications are unaffected for a minimum of three days.
- Unless there is a good reason for using a third repository, such as using a geographically local full repository for geographical reasons, use the two repository design. Having three full repositories means that you never know which are the two that are currently in use, and there might be administrative problems caused by interactions between multiple workload management parameters. It is not recommend to have more than two full repositories.
- If you still need better availability, consider hosting the full repository queue managers as multi-instance queue managers or using platform specific high availability support to improve their availability.
- You must fully interconnect all the full repository queue managers with manually defined cluster sender channels. Particular care must be taken when the cluster does have, for some justifiable reason, more than two full repositories. In this situation it is often possible to miss one or more channels and for it not to be immediately apparent. When full interconnection does not occur, hard to diagnose problems often arise. They are hard to diagnose because some full repositories not holding all repository data and therefore resulting in queue managers in the cluster having different views of the cluster depending on the full repositories that they connect to.

Should applications use queues on full repositories?

A full repository is in most ways exactly like any other queue manager, and it is therefore possible to host application queues on the full repository and connect applications directly to these queue managers. Should applications use queues on full repositories?

The commonly accepted answer is "No?". Although this configuration is possible, many customers prefer to keep these queue managers dedicated to maintaining the full repository cluster cache. Points to consider when deciding on either option are described here, but ultimately the cluster architecture must be appropriate to the particular demands of the environment.

- **Upgrades:** Usually, in order to use new cluster features in new releases of IBM MQ the full repository queue managers of that cluster must be upgraded first. When an application in the cluster wants to use new features, it might be useful to be able to update the full repositories (and some subset of partial repositories) without testing a number of co-located applications.
- **Maintenance:** In a similar way if you must apply urgent maintenance to the full repositories, they can be restarted or refreshed with the **REFRESH** command without touching applications.
- **Performance:** As clusters grow and demands on the full repository cluster cache maintenance become greater, keeping applications separate reduces risk of this affecting application performance through contention for system resources.
- **Hardware requirements:** Typically, full repositories do not need to be powerful; for example, a simple UNIX server with a good expectation of availability is sufficient. Alternatively, for very large or constantly changing clusters, the performance of the full repository computer must be considered.
- **Software requirements:** Requirements are usually the main reason for choosing to host application queues on a full repository. In a small cluster, collocation might mean a requirement for fewer queue managers/servers over all.

Managing channel definitions

Even within a single cluster, multiple channel definitions can exist giving multiple routes between two queue managers.

There is sometimes an advantage to having parallel channels within a single cluster, but this design decision must be considered thoroughly; apart from adding complexity, this design might result in channels being under-utilized which reduces performance. This situation occurs because testing usually involves sending lots of messages at a constant rate, so the parallel channels are fully used. But with real-world conditions of a non-constant stream of messages, the workload balancing algorithm causes performance to drop as the message flow is switched from channel to channel.

When a queue manager is a member of multiple clusters, the option exists to use a single channel definition with a cluster namelist, rather than defining a separate CLUSRCVR channel for each cluster. However, this setup can cause administration difficulties later; consider for example the case where SSL is to be applied to one cluster but not a second. It is therefore preferable to create separate definitions, and the naming convention suggested in "Cluster naming conventions" on page 29 supports this.

Workload balancing over multiple channels

This information is intended as an advanced understanding of the subject. For the basic explanation of this subject (which must be understood before using the information here), see Using clusters for workload management, Workload balancing in clusters, and The cluster workload management algorithm.

The cluster workload management algorithm provides a large set of tools, but they must not all be used with each other without fully understanding how they work and interact. It might not be immediately obvious how important channels are to the workload balancing process: The workload management

round-robin algorithm behaves as though multiple cluster channels to a queue manager that owns a clustered queue, are treated as multiple instances of that queue. This process is explained in more detail in the following example:

1. There are two queue managers hosting a queue in a cluster: QM1 and QM2.
2. There are five cluster receiver channels to QM1.
3. There is only one cluster receiver channel to QM2.
4. When **MQPUT** or **MQOPEN** on QM3 chooses an instance, the algorithm is five times more likely to send the message to QM1 than to QM2.
5. The situation in step 4 occurs because the algorithm sees six options to choose from (5+1) and round-robins across all five channels to QM1 and the single channel to QM2.

Another subtle behavior is that even when putting messages to a clustered queue that happens to have one instance configured on the local queue manager, IBM MQ uses the state of the local cluster receiver channel to decide whether messages are to be put to the local instance of the queue or remote instances of the queue. In this scenario:

1. When putting messages the workload management algorithm does not look at individual cluster queues, it looks at the cluster channels which can reach those destinations.
2. To reach local destinations, the local receiver channels are included in this list (although they are not used to send the message).
3. When a local receiver channel is stopped, the workload management algorithm, prefers an alternative instance by default if its CLUSRCVR is not stopped. If there are multiple local CLUSRCVR instances for the destination and at least one is not stopped, the local instance remains eligible.

Clustering: Application isolation using multiple cluster transmission queues:

You can isolate the message flows between queue managers in a cluster. You can place messages being transported by different cluster-sender channels onto different cluster transmission queues. You can use the approach in a single cluster or with overlapping clusters. The topic provides examples and some best practices to guide you in choosing an approach to use.

When you deploy an application, you have a choice over which IBM MQ resources it shares with other applications and which resources it does not share. There are a number of types of resources that can be shared, the main ones being the server itself, the queue manager, channels, and queues. You can choose to configure applications with fewer shared resources; allocating separate queues, channels, queue managers, or even servers to individual applications. If you do so, the overall system configuration becomes bigger and more complex. Using IBM MQ clusters reduces the complexity of managing more servers, queue managers, queues, and channels, but it introduces another shared resource, the cluster transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

Figure 12 on page 36 is a slice out of a large IBM MQ deployment that illustrates the significance of sharing `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. In the diagram, the application, `Client App`, is connected to the queue manager QM2 in cluster CL1. A message from `Client App` is processed by the application, `Server App`. The message is retrieved by `Server App` from the cluster queue Q1 on the queue manager QM3 in CLUSTER2. Because the client and server applications are not in the same cluster, the message is transferred by the gateway queue manager QM1.

The normal way to configure a cluster gateway is to make the gateway queue manager a member of all the clusters. On the gateway queue manager are defined clustered alias queues for cluster queues in all the clusters. The clustered queue aliases are available in all the clusters. Messages put to the cluster queue aliases are routed via the gateway queue manager to their correct destination. The gateway queue manager puts messages sent to the clustered alias queues onto the common `SYSTEM.CLUSTER.TRANSMIT.QUEUE` on QM1.

The hub and spoke architecture requires all messages between clusters to pass through the gateway queue manager. The result is that all messages flow through the single cluster transmission queue on QM1, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

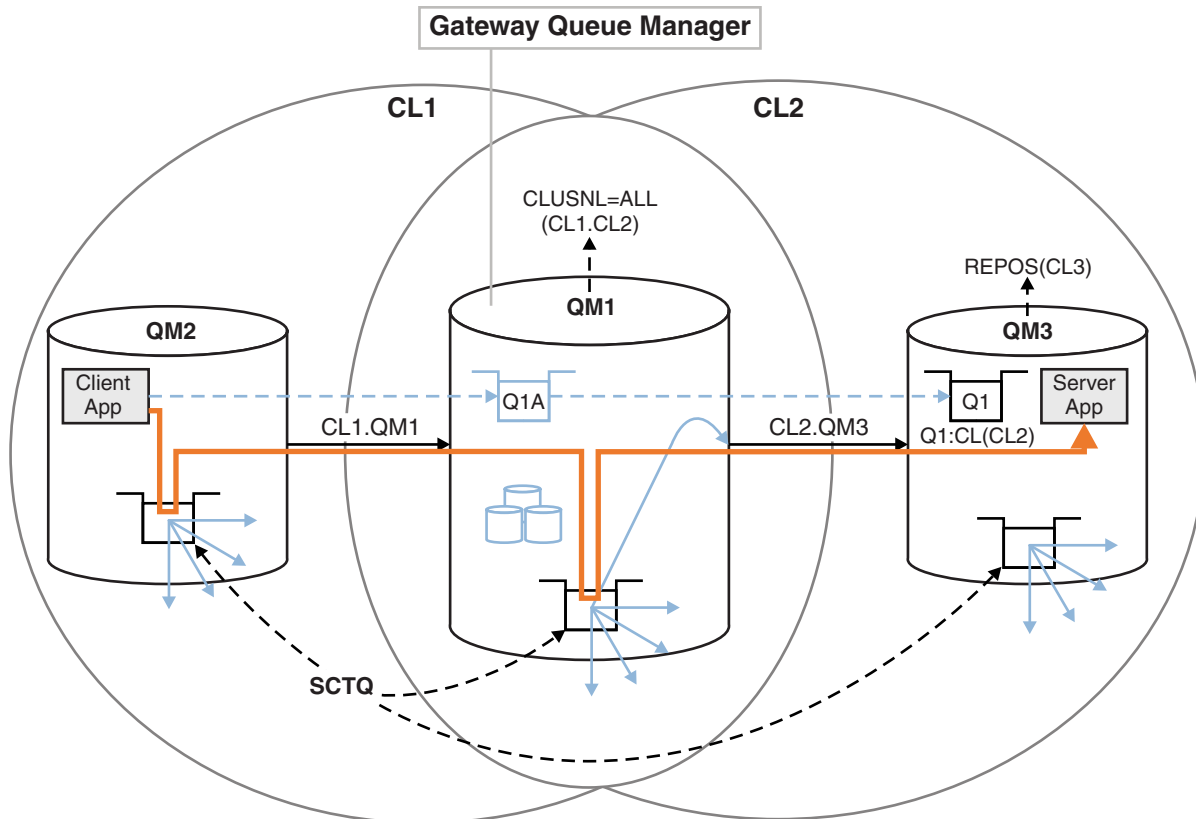
From a performance perspective, a single queue is not a problem. A common transmission queue generally does not represent a performance bottleneck. Message throughput on the gateway is largely determined by the performance of the channels that connect to it. Throughput is not generally affected by the number of queues, or the number of messages on the queues that use the channels.

From some other perspectives, using a single transmission queue for multiple applications has drawbacks:

- You cannot isolate the flow of messages to one destination from the flow of messages to another destination. You cannot separate the storage of messages before they are forwarded, even if the destinations are in different clusters on different queue managers.

If one cluster destination becomes unavailable, messages for that destination build-up in the single transmission queue, and eventually the messages fill it up. Once the transmission queue is full, it stops messages from being placed onto the transmission queue for any cluster destination.

- It is not easy to monitor the transfer of messages to different cluster destinations. All the messages are on the single transmission queue. Displaying the depth of the transmission queue gives you little indication whether messages are being transferred to all destinations.



Note: The arrows in Figure 12 and following figures are of different types. Solid arrows represent message flows. The labels on solid arrows are message channel names. The gray solid arrows are potential message flows from the SYSTEM.CLUSTER.TRANSMIT.QUEUE onto cluster-sender channels. Black dashed lines connect labels to their targets. Gray dashed arrows are references; for example from an MQOPEN call by Client App to the cluster alias queue definition Q1A.

Figure 12. Client-server application deployed to hub and spoke architecture using IBM MQ clusters

In Figure 12, clients of Server App open the queue Q1A. Messages are put to SYSTEM.CLUSTER.TRANSMIT.QUEUE on QM2, transferred to SYSTEM.CLUSTER.TRANSMIT.QUEUE on QM1, and then transferred to Q1 on QM3, where they are received by the Server App application.

The message from Client App passes through system cluster transmission queues on QM2 and QM1. In Figure 12, the objective is to isolate the message flow on the gateway queue manager from the client application, so that its messages are not stored on SYSTEM.CLUSTER.TRANSMIT.QUEUE. You can isolate flows on any of the other clustered queue managers. You can also isolate flows in the other direction, back to the client. To keep the descriptions of the solutions brief, the descriptions consider only a single flow from the client application.

Solutions to isolating cluster message traffic on a cluster gateway queue manager

One way to solve the problem is to use queue manager aliases, or remote queue definitions, to bridge between clusters. Create a clustered remote queue definition, a transmission queue, and a channel, to separate each message flow on the gateway queue manager; see Adding a remote queue definition to isolate messages sent from a gateway queue manager.

From Version 7.5 onwards, cluster queue managers are not limited to a single cluster transmission queue. You have two choices:

1. Define additional cluster transmission queues manually, and define which cluster-sender channels transfer messages from each transmission queue; see Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager.
2. Allow the queue manager to create and manage additional cluster transmission queues automatically. It defines a different cluster transmission queue for each cluster-sender channel; see Changing the default to separate cluster transmission queues to isolate message traffic.

You can combine manually defined cluster transmission queues for some cluster-sender channels, with the queue manager managing the rest. The combination of transmission queues is the approach taken in Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager. In that solution, most messages between clusters use the common `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. One application is critical, and all its message flows are isolated from other flows by using one manually defined cluster transmission queue.

The configuration in Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager is limited. It does not separate the message traffic going to a cluster queue on the same queue manager in the same cluster as another cluster queue. You can separate the message traffic to individual queues by using the remote queue definitions that are part of distributed queuing. With clusters, using multiple cluster transmission queues, you can separate message traffic that goes to different cluster-sender channels. Multiple cluster queues in the same cluster, on the same queue manager, share a cluster-sender channel. Messages for those queues are stored on the same transmission queue, before being forwarded from the gateway queue manager. In the configuration in Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager, the limitation is side-stepped by adding another cluster and making the queue manager and cluster queue a member of the new cluster. The new queue manager might be the only queue manager in the cluster. You could add more queue managers to the cluster, and use the same cluster to isolate cluster queues on those queue managers as well.

Related concepts:

“Access control and multiple cluster transmission queues” on page 13

Choose between three modes of checking when an application puts messages to remote cluster queues.

The modes are checking remotely against the cluster queue, checking locally against `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, or checking against local profiles for the cluster queue, or cluster queue manager.

“Overlapping clusters” on page 20

Overlapping clusters provide additional administrative capabilities. Use namelists to reduce the number of commands needed to administer overlapping clusters.

Related information:

Authorizing putting messages on remote cluster queues

Working with cluster transmission queues and cluster-sender channels

Adding a remote queue definition to isolate messages sent from a gateway queue manager

Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager

Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager

Changing the default to separate cluster transmission queues to isolate message traffic

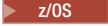
Creating two-overlapping clusters with a gateway queue manager

Configuring message paths between clusters

Security


setmqaut

Clustering: Planning how to configure cluster transmission queues:

You are guided through the choices of cluster transmission queues. You can configure one common default queue, separate default queues, or manually defined queues.  To configure a cluster-sender channel to use a transmission queue other than `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, you need to enable version 8 new function, using the mode of operation (OPMODE) system parameter in the `CSQ6SYSP` parameter.

Before you begin

Review “How to choose what type of cluster transmission queue to use” on page 40.

 See the mode of operation (OPMODE) topic for more information.

About this task

You have some choices to make when you are planning how to configure a queue manager to select a cluster transmission queue.

1. What is the default cluster transmission queue for cluster message transfers?
 - a. A common cluster transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.
 - b. Separate cluster transmission queues. The queue manager manages the separate cluster transmission queues. It creates them as permanent-dynamic queues from the model queue, `SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE`. It creates one cluster transmission queue for each cluster-sender channel it uses.
2. For the cluster transmission queues that you do decide to create manually, you have a further two choices:
 - a. Define a separate transmission queue for each cluster-sender channel that you decide to configure manually. In this case, set the **CLCHNAME** queue attribute of the transmission queue to the name of a cluster-sender channel. Select the cluster-sender channel that is to transfer messages from this transmission queue.
 - b. Combine message traffic for a group of cluster-sender channels onto the same cluster transmission queue; see Figure 13 on page 39. In this case, set the **CLCHNAME** queue attribute of each common transmission queue to a generic cluster-sender channel name. A generic cluster-sender channel name is a filter to group cluster-sender channel names. For example, `SALES.*` groups all cluster-sender channels that have names beginning with `SALES.`. You can place multiple wildcard characters anywhere in the filter-string. The wildcard character is an asterisk, “*”. It represents from zero to any number of characters.

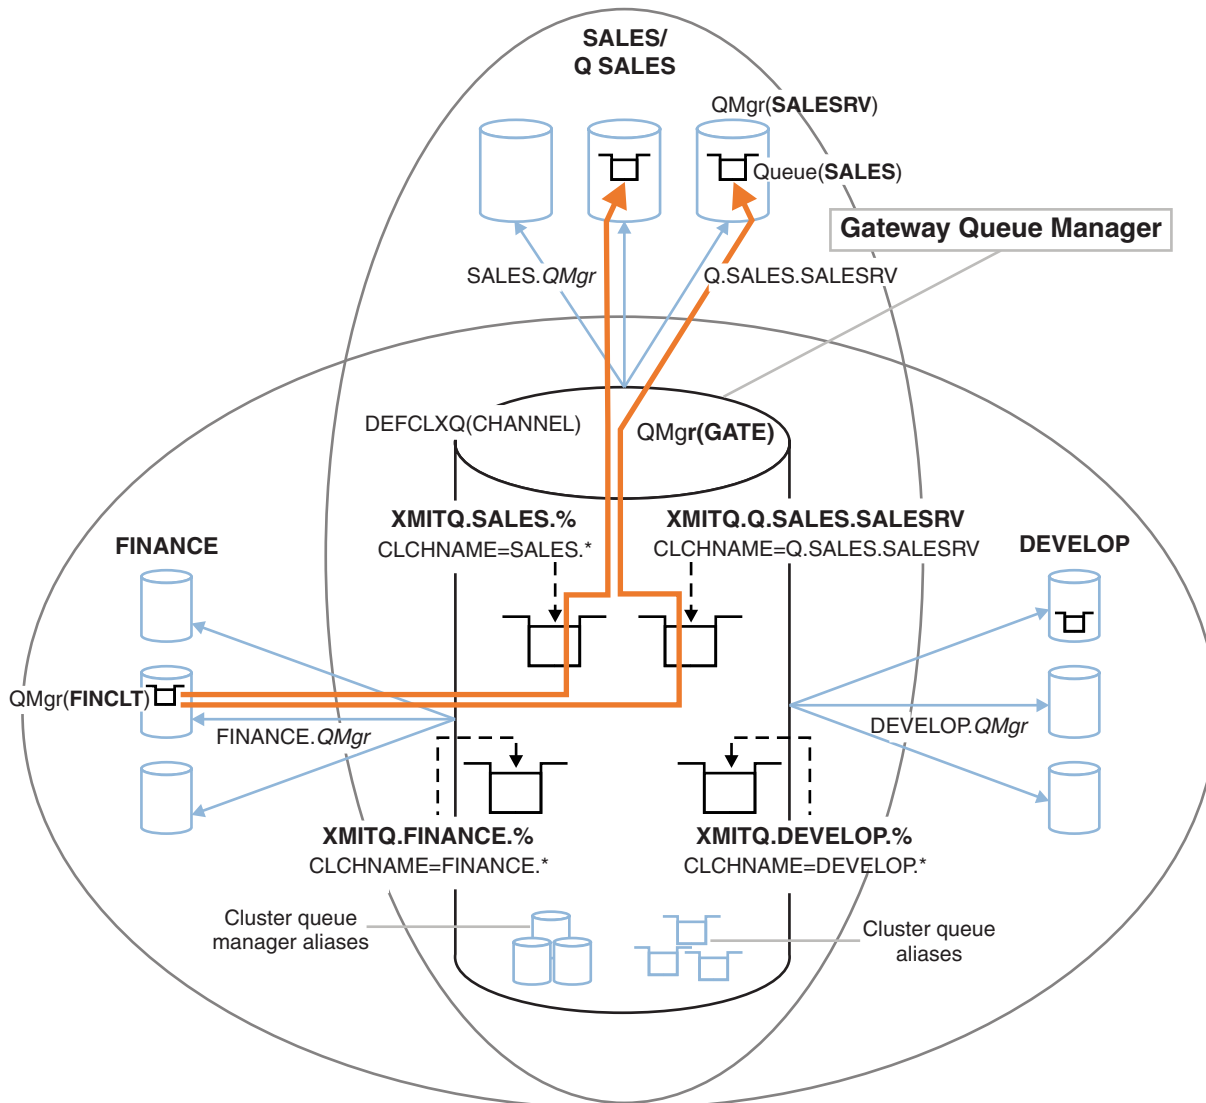


Figure 13. Example of specific transmission queues for different departmental IBM MQ clusters

Procedure

1. Select the type of default cluster transmission queue to use.
 - Choose a single cluster transmission queue, or separate queues for each cluster connection.

Leave the default setting or run the **MQSC** command:

```
ALTER QMGR DEFCLXQ(CHANNEL)
```
2. Isolate any message flows that must not share a cluster transmission queue with other flows.
 - See “Clustering: Example configuration of multiple cluster transmission queues” on page 42. In the example the SALES queue, which must be isolated, is a member of the SALES cluster, on SALESRV. To isolate the SALES queue, create a new cluster Q.SALES, make the SALESRV queue manager a member, and modify the SALES queue to belong to Q.SALES.
 - Queue managers that send messages to SALES must also be members of the new cluster. If you use a clustered queue alias and a gateway queue manager, as in the example, in many cases you can limit the changes to making the gateway queue manager a member of the new cluster.
 - However, separating flows from the gateway to the destination does not separate flows to the gateway from the source queue manager. But it sometimes turns out to be sufficient to separate flows from the gateway and not flows to the gateway. If it is not sufficient, then add the source

queue manager into the new cluster. If you want messages to travel through the gateway, move the cluster alias to the new cluster and continue to send messages to the cluster alias on the gateway, and not directly to the target queue manager.

Follow these steps to isolate message flows:

- a. Configure the destinations of the flows so that each target queue is the only queue in a specific cluster, on that queue manager.
 - b. Create the cluster-sender and cluster-receiver channels for any new clusters you have created following a systematic naming convention.
 - See “Clustering: Special considerations for overlapping clusters” on page 28.
 - c. Define a cluster transmission queue for each isolated destination on every queue manager that sends messages to the target queue.
 - A naming convention for cluster transmission queues is to use the value of the cluster channel name attribute, CLCHNAME, prefixed with XMITQ.
3. Create cluster transmission queues to meet governance or monitoring requirements.
- Typical governance and monitoring requirements result in a transmission queue per cluster or a transmission queue per queue manager. If you follow the naming convention for cluster channels, *ClusterName.QueueManagerName*, it is easy to create generic channel names that select a cluster of queue managers, or all the clusters a queue manager is a member of; see “Clustering: Example configuration of multiple cluster transmission queues” on page 42.
 - Extend the naming convention for cluster transmission queues to cater for generic channel names, by replacing the asterisk symbol by a percent sign. For example,
`DEFINE QLOCAL(XMITQ.SALES.%) USAGE(XMITQ) CLCHNAME(SALES.*)`

Related concepts:

“Access control and multiple cluster transmission queues” on page 13

Choose between three modes of checking when an application puts messages to remote cluster queues. The modes are checking remotely against the cluster queue, checking locally against `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, or checking against local profiles for the cluster queue, or cluster queue manager.

“Overlapping clusters” on page 20

Overlapping clusters provide additional administrative capabilities. Use namelists to reduce the number of commands needed to administer overlapping clusters.

Related information:

Working with cluster transmission queues and cluster-sender channels

Adding a remote queue definition to isolate messages sent from a gateway queue manager

Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager

Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager

Changing the default to separate cluster transmission queues to isolate message traffic

Creating two-overlapping clusters with a gateway queue manager

Configuring message paths between clusters

How to choose what type of cluster transmission queue to use:

How to choose between different cluster transmission queue configuration options.

From Version 7.5 onwards, you can choose which cluster transmission queue is associated with a cluster-sender channel.

1. You can have all cluster-sender channels associated with the single default cluster transmit queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. This option is the default, and is the only choice for queue managers that run version Version 7.1, or earlier.

2. You can set all cluster-sender channels to be automatically associated with a separate cluster transmission queue. The queues are created by the queue manager from the model queue `SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE` and named `SYSTEM.CLUSTER.TRANSMIT.ChannelName`. Channels will use their uniquely named cluster transmit queue if the queue manager attribute `DEFCLXQ` is set to `CHANNEL`.
3. You can set specific cluster-sender channels to be served by a single cluster transmission queue. Select this option by creating a transmission queue and setting its `CLCHNAME` attribute to the name of the cluster-sender channel.
4. You can select groups of cluster-sender channels to be served by a single cluster transmission queue. Select this option by creating a transmission queue and setting its `CLCHNAME` attribute to a generic channel name, such as `ClusterName.*`. If you name cluster channels by following the naming conventions in “Clustering: Special considerations for overlapping clusters” on page 28, this name selects all cluster channels connected to queue managers in the cluster `ClusterName`.

You can combine either of the default cluster transmission queue options for some cluster-sender channels, with any number of specific and generic cluster transmission queue configurations.

Best practices

In most cases, for existing IBM MQ installations, the default configuration is the best choice. A cluster queue manager stores cluster messages on a single cluster transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. You have the choice of changing the default to storing messages for different queue managers and different clusters on separate transmission queues, or of defining your own transmission queues.

In most cases, for new IBM MQ installations, the default configuration is also the best choice. The process of switching from the default configuration to the alternative default of having one transmission queue for each cluster-sender channel is automatic. Switching back is also automatic. The choice of one or the other is not critical, you can reverse it.

The reason for choosing a different configuration is more to do with governance, and management, than with functionality or performance. With a couple of exceptions, configuring multiple cluster transmission queues does not benefit the behavior of the queue manager. It results in more queues, and requires you to modify monitoring and management procedures you have already set up that refer to the single transmission queue. That is why, on balance, remaining with the default configuration is the best choice, unless you have strong governance or management reasons for a different choice.

The exceptions are both concerned with what happens if the number of messages stored on `SYSTEM.CLUSTER.TRANSMIT.QUEUE` increases. If you take every step to separate the messages for one destination from the messages for another destination, then channel and delivery problems with one destination ought not to affect the delivery to another destination. However, the number of messages stored on `SYSTEM.CLUSTER.TRANSMIT.QUEUE` can increase due to not delivering messages fast enough to one destination. The number of messages on `SYSTEM.CLUSTER.TRANSMIT.QUEUE` for one destination can affect the delivery of messages to other destinations.

To avoid problems that result from filling up a single transmission queue, aim to build sufficient capacity into your configuration. Then, if a destination fails and a message backlog starts to build up, you have time to fix the problem.

If messages are routed through a hub queue manager, such as a cluster gateway, they share a common transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. If the number of messages stored on `SYSTEM.CLUSTER.TRANSMIT.QUEUE` on the gateway queue manager reaches its maximum depth, the queue manager starts to reject new messages for the transmission queue until its depth reduces. The congestion affects messages for all destinations that are routed through the gateway. Messages back up the transmission queues of other queue managers that are sending messages to the gateway. The problem

manifests itself in messages written to queue manager error logs, falling message throughput, and longer elapsed times between sending a message and the time that a message arrives at its destination.

The effect of congestion on a single transmission queue can become apparent, even before it is full. If you have a mixed message traffic, with some large non-persistent messages and some small messages, the time to deliver small messages increases as the transmission queue fills. The delay is due to writing large non-persistent messages to disk that would not normally get written to disk. If you have time critical message flows, sharing a cluster transmission queue with other mixed messages flows, it could be worth configuring a special message path to isolate it from other message flows; see Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager.

The other reasons for configuring separate cluster transmission queues are to meet governance requirements, or to simplify monitoring messages that are sent to different cluster destinations. For example, you might have to demonstrate that messages for one destination never share a transmission queue with messages for another destination.

Change the queue manager attribute **DEFCLXQ** that controls the default cluster transmission queue, to create different cluster transmission queues for every cluster-sender channel. Multiple destinations can share a cluster-sender channel, so you must plan your clusters to meet this objective fully. Apply the method Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager systematically to all your cluster queues. The result you are aiming for is for no cluster destination to share a cluster-sender channel with another cluster destination. As a consequence, no message for a cluster destination shares its cluster transmission queue with a message for another destination.

Creating a separate cluster transmission queue for some specific message flow, makes it easy to monitor the flow of messages to that destination. To use a new cluster transmission queue, define the queue, associate it with a cluster-sender channel, and stop and start the channel. The change does not have to be permanent. You could isolate a message flow for a while, to monitor the transmission queue, and then revert to using the default transmission queue again.

Clustering: Example configuration of multiple cluster transmission queues:

In this task you apply the steps to plan multiple cluster transmission queues to three overlapping clusters. The requirements are to separate messages flows to one cluster queue, from all other message flows, and to store messages for different clusters on different cluster transmission queues.

About this task

The steps in this task show how to apply the procedure in “Clustering: Planning how to configure cluster transmission queues” on page 38 and arrive at the configuration shown in Figure 14 on page 43. It is an example of three overlapping clusters, with a gateway queue manager, that is configured with separate cluster transmission queues. The MQSC commands to define the clusters are described in “Creating the example clusters” on page 45.

For the example, there are two requirements. One is to separate the message flow from the gateway queue manager to the sales application that logs sales. The second is to query how many messages are waiting to be sent to different departmental areas at any point in time. The SALES, FINANCE, and DEVELOP clusters are already defined. Cluster messages are currently forwarded from `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

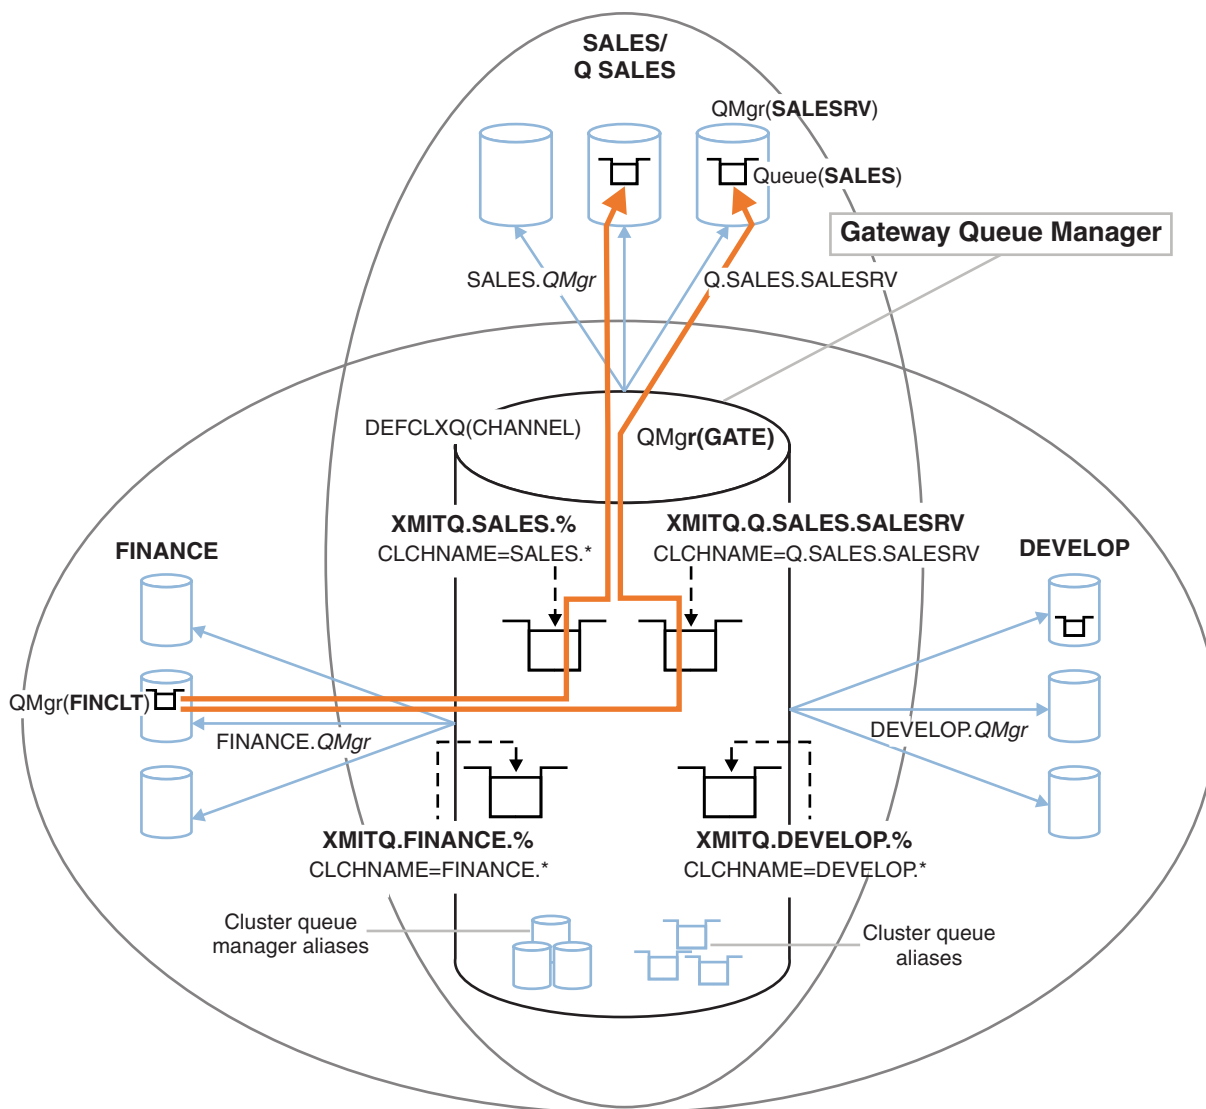


Figure 14. Example of specific transmission queues for different departmental IBM MQ clusters

The steps to modify the clusters are as follows; see Figure 16 on page 47 for the definitions.

Procedure

1. The first configuration step is to “ Select the type of default cluster transmission queue to use “. The decision is to create separate default cluster transmission queues by running the following **MQSC** command on the GATE queue manager.

```
ALTER QMGR DEFCLXQ(CHANNEL)
```

There is no strong reason for choosing this default, as the intention is to manually define cluster transmission queues. The choice does have a weak diagnostic value. If a manual definition is done wrongly, and a message flows down a default cluster transmission queue, it shows up in the creation of a permanent-dynamic cluster transmission queue.

2. The second configuration step is to “ Isolate any message flows that must not share a cluster transmission queue with other flows “. In this case the sales application that receives messages from the queue SALES on SALESRV requires isolation. Only isolation of messages from the gateway queue manager is required. The three sub-steps are:

- a. “ Configure the destinations of the flows so that each target queue is the only queue in a specific cluster, on that queue manager ”.

The example requires adding the queue manager SALESRV to a new cluster within the sales department. If you have few queues that require isolation, you might decide on creating a specific cluster for the SALES queue. A possible naming convention for the cluster name is to name such clusters, *Q.QueueName*, for example Q.SALES. An alternative approach, which might be more practical if you have a large number of queues to be isolated, is to create clusters of isolated queues where and when needed. The clusters names might be QUEUES. *n*.

In the example, the new cluster is called Q.SALES. To add the new cluster, see the definitions in Figure 16 on page 47. The summary of definition changes is as follows:

- 1) Add Q.SALES to the namelist of clusters on the repository queue managers. The namelist is referred to in the queue manager **REPOSNL** parameter.
- 2) Add Q.SALES to the namelist of clusters on the gateway queue manager. The namelist is referred to in all the cluster queue alias and cluster queue manager alias definitions on the gateway queue manager.
- 3) Create a namelist on the queue manager SALESRV, for both the clusters it is a member of, and change the cluster membership of the SALES queue:

```
DEFINE NAMLIST(CLUSTERS) NAMES(SALES, Q.SALES) REPLACE
ALTER QLOCAL(SALES) CLUSTER(' ') CLUSNL(SALESRV.CLUSTERS)
```

The SALES queue is a member of both clusters, just for the transition. Once the new configuration is running, you remove the SALES queue from the SALES cluster; see Figure 17 on page 47.

- b. “ Create the cluster-sender and cluster-receiver channels for any new clusters you have created following a systematic naming convention ”.
 - 1) Add the cluster-receiver channel Q.SALES. *RepositoryQMgr* to each of the repository queue managers
 - 2) Add the cluster-sender channel Q.SALES. *OtherRepositoryQMgr* to each of the repository queue managers, to connect to the other repository manager. Start these channels.
 - 3) Add the cluster receiver channels Q.SALES.SALESRV, and Q.SALES.GATE to either of the repository queue managers that is running.
 - 4) Add the cluster-sender channels Q.SALES.SALESRV, and Q.SALES.GATE to the SALESRV and GATE queue managers. Connect the cluster-sender channel to the repository queue manager that you created the cluster-receiver channels on.
- c. “ Define a cluster transmission queue for each isolated destination on every queue manager that sends messages to the target queue ”.

On the gateway queue manager define the cluster transmission queue XMITQ.Q.SALES.SALESRV for the Q.SALES.SALESRV cluster-sender channel:

```
DEFINE QLOCAL(XMITQ.Q.SALES.SALESRV) USAGE(XMITQ) CLCHNAME(Q.SALES.SALESRV) REPLACE
```

3. The third configuration step is to “ Create cluster transmission queues to meet governance or monitoring requirements ”.

On the gateway queue manager define the cluster transmission queues:

```
DEFINE QLOCAL(XMITQ.SALES) USAGE(XMITQ) CLCHNAME(SALES.*) REPLACE
DEFINE QLOCAL(XMITQ.DEVELOP) USAGE(XMITQ) CLCHNAME(DEVELOP.*) REPLACE
DEFINE QLOCAL(XMITQ.FINANCE) USAGE(XMITQ) CLCHNAME(SALES.*) REPLACE
```

What to do next

Switch to the new configuration on the gateway queue manager.

The switch is triggered by starting the new channels, and restarting the channels that are now associated with different transmission queues. Alternatively, you can stop and start the gateway queue manager.

1. Stop the following channels on the gateway queue manager:

```
SALES. Qmgr
DEVELOP. Qmgr
FINANCE. Qmgr
```

2. Start the following channels on the gateway queue manager:

```
SALES. Qmgr
DEVELOP. Qmgr
FINANCE. Qmgr
Q.SALES.SAVESRV
```

When the switch is complete, remove the SALES queue from the SALES cluster; see Figure 17 on page 47.

Creating the example clusters:

The definitions and instructions to create the example cluster, and modify it to isolate the SALES queue and separate messages on the gateway queue manager.

About this task

The full **MQSC** commands to create the FINANCE, SALES, and Q.SALES clusters are provided in Figure 15 on page 46, Figure 16 on page 47, and Figure 17 on page 47. The definitions for the basic clusters are in Figure 15 on page 46. The definitions in Figure 16 on page 47 modify the basic clusters to isolate the SALES queue, and to separate cluster messages to FINANCE and SALES. Finally, to remove the SALES queue from the SALES cluster; see Figure 17 on page 47. The DEVELOP cluster is omitted from the definitions, to keep the definitions shorter.

Procedure

1. Create the SALES and FINANCE clusters, and the gateway queue manager.

- a. Create the queue managers.

Run the command: `crtmqm -sax -u SYSTEM.DEAD.LETTER.QUEUE QmgrName` for each of the queue manager names in Table 4.

Table 4. Queue manager names and port numbers

Description	Queue Manager Name	Port number
Finance repository	FINR1	1414
Finance repository	FINR2	1415
Finance client	FINCLT	1418
Sales repository	SALER1	1416
Sales repository	SALER2	1417
Sales server	SALESRV	1419
Gateway	GATE	1420

- b. Start all the queue managers

Run the command: `strmqm QmgrName` for each of the queue manager names in Table 4.

- c. Create the definitions for each of the queue managers

Run the command: `runmqsc QmgrName < filename` where the files are listed in Figure 15 on page 46, and the file name matches the queue manager name.

finr1.txt

```

DEFINE LISTENER(1414) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1414) REPLACE
START LISTENER(1414)
ALTER QMGR REPOS(FINANCE)
DEFINE CHANNEL(FINANCE.FINR2) CHLTYPE(CLUSSDR) CONNAME('localhost(1415)') CLUSTER(FINANCE) REPLACE
DEFINE CHANNEL(FINANCE.FINR1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1414)') CLUSTER(FINANCE) REPLACE

```

finr2.txt

```

DEFINE LISTENER(1415) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1415) REPLACE
START LISTENER(1415)
ALTER QMGR REPOS(FINANCE)
DEFINE CHANNEL(FINANCE.FINR1) CHLTYPE(CLUSSDR) CONNAME('localhost(1414)') CLUSTER(FINANCE) REPLACE
DEFINE CHANNEL(FINANCE.FINR2) CHLTYPE(CLUSRCVR) CONNAME('localhost(1415)') CLUSTER(FINANCE) REPLACE

```

finclt.txt

```

DEFINE LISTENER(1418) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1418) REPLACE
START LISTENER(1418)
DEFINE CHANNEL(FINANCE.FINR1) CHLTYPE(CLUSSDR) CONNAME('localhost(1414)') CLUSTER(FINANCE) REPLACE
DEFINE CHANNEL(FINANCE.FINCLT) CHLTYPE(CLUSRCVR) CONNAME('localhost(1418)') CLUSTER(FINANCE) REPLACE
DEFINE QMODEL(SYSTEM.SAMPLE.REPLY) REPLACE

```

saler1.txt

```

DEFINE LISTENER(1416) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1416) REPLACE
START LISTENER(1416)
ALTER QMGR REPOS(SALES)
DEFINE CHANNEL(SALES.SALER2) CHLTYPE(CLUSSDR) CONNAME('localhost(1417)') CLUSTER(SALES) REPLACE
DEFINE CHANNEL(SALES.SALER1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1416)') CLUSTER(SALES) REPLACE

```

saler2.txt

```

DEFINE LISTENER(1417) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1417) REPLACE
START LISTENER(1417)
ALTER QMGR REPOS(SALES)
DEFINE CHANNEL(SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('localhost(1416)') CLUSTER(SALES) REPLACE
DEFINE CHANNEL(SALES.SALER2) CHLTYPE(CLUSRCVR) CONNAME('localhost(1417)') CLUSTER(SALES) REPLACE

```

salesrv.txt

```

DEFINE LISTENER(1419) TRPTYPE(TCP) IPADDR(localhost) CONTROL(QMGR) PORT(1419) REPLACE
START LISTENER(1419)
DEFINE CHANNEL(SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('localhost(1416)') CLUSTER(SALES) REPLACE
DEFINE CHANNEL(SALES.SALESRV) CHLTYPE(CLUSRCVR) CONNAME('localhost(1419)') CLUSTER(SALES) REPLACE
DEFINE QLOCAL(SALES) CLUSTER(SALES) TRIGGER INITQ(SYSTEM.DEFAULT.INITIATION.QUEUE) PROCESS(ECHO) REPLACE
DEFINE PROCESS(ECHO) APPLICID(AMQSECH) REPLACE

```

gate.txt

```

DEFINE LISTENER(1420) TRPTYPE(TCP) IPADDR(LOCALHOST) CONTROL(QMGR) PORT(1420) REPLACE
START LISTENER(1420)
DEFINE NAMELIST(ALL) NAMES(SALES, FINANCE)
DEFINE CHANNEL(FINANCE.FINR1) CHLTYPE(CLUSSDR) CONNAME('LOCALHOST(1414)') CLUSTER(FINANCE) REPLACE
DEFINE CHANNEL(FINANCE.GATE) CHLTYPE(CLUSRCVR) CONNAME('LOCALHOST(1420)') CLUSTER(FINANCE) REPLACE
DEFINE CHANNEL(SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('LOCALHOST(1416)') CLUSTER(SALES) REPLACE
DEFINE CHANNEL(SALES.GATE) CHLTYPE(CLUSRCVR) CONNAME('LOCALHOST(1420)') CLUSTER(SALES) REPLACE
DEFINE QALIAS(A.SALES) CLUSNL(ALL) TARGET(SALES) TARGTYPE(QUEUE) DEFBIND(NOTFIXED) REPLACE
DEFINE QREMOTE(FINCLT) RNAME(' ') RQMNAME(FINCLT) CLUSNL(ALL) REPLACE
DEFINE QREMOTE(SALESRV) RNAME(' ') RQMNAME(SALESRV) CLUSNL(ALL) REPLACE

```

Figure 15. Definitions for the basic clusters

2. Test the configuration by running the sample request program.
 - a. Start the trigger monitor program on the SALESRV queue manager

On Windows, open a command window and run the command `runmqtrm -m SALESRV`
 - b. Run the sample request program, and send a request.

On Windows, open a command window and run the command `amqsreq A.SALES FINCLT`

The request message is echoed back, and after 15 seconds the sample program finishes.

3. Create the definitions to isolate the SALES queue in the Q.SALES cluster and separate cluster messages for the SALES and FINANCE cluster on the gateway queue manager.

Run the command: `runmqsc QmgrName < filename` where the files are listed in Figure 16, and the file name almost matches the queue manager name.

chgsaler1.txt

```
DEFINE NAMLIST(CLUSTERS) NAMES(SALES, Q.SALES)
ALTER QMGR REPOS(' ') REPOSNL(CLUSTERS)
DEFINE CHANNEL(Q.SALES.SALER2) CHLTYPE(CLUSSDR) CONNAME('localhost(1417)') CLUSTER(Q.SALES) REPLACE
DEFINE CHANNEL(Q.SALES.SALER1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1416)') CLUSTER(Q.SALES) REPLACE
```

chgsaler2.txt

```
DEFINE NAMLIST(CLUSTERS) NAMES(SALES, Q.SALES)
ALTER QMGR REPOS(' ') REPOSNL(CLUSTERS)
DEFINE CHANNEL(Q.SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('localhost(1416)') CLUSTER(Q.SALES) REPLACE
DEFINE CHANNEL(Q.SALES.SALER2) CHLTYPE(CLUSRCVR) CONNAME('localhost(1417)') CLUSTER(Q.SALES) REPLACE
```

chgsalesrv.txt

```
DEFINE NAMLIST (CLUSTERS) NAMES(SALES, Q.SALES)
DEFINE CHANNEL(Q.SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('localhost(1416)') CLUSTER(Q.SALES) REPLACE
DEFINE CHANNEL(Q.SALES.SAVESRV) CHLTYPE(CLUSRCVR) CONNAME('localhost(1419)') CLUSTER(Q.SALES) REPLACE
ALTER QLOCAL (SALES) CLUSTER(' ') CLUSNL(CLUSTERS)
```

chggate.txt

```
ALTER NAMLIST(ALL) NAMES(SALES, FINANCE, Q.SALES)
ALTER QMGR DEFCLXQ(CHANNEL)
DEFINE CHANNEL(Q.SALES.SALER1) CHLTYPE(CLUSSDR) CONNAME('localhost(1416)') CLUSTER(Q.SALES) REPLACE
DEFINE CHANNEL(Q.SALES.GATE) CHLTYPE(CLUSRCVR) CONNAME('localhost(1420)') CLUSTER(Q.SALES) REPLACE
DEFINE QLOCAL (XMITQ.Q.SALES.SALESRV) USAGE(XMITQ) CLCHNAME(Q.SALES.SALESRV) REPLACE
DEFINE QLOCAL (XMITQ.SALES) USAGE(XMITQ) CLCHNAME(SALES.*) REPLACE
DEFINE QLOCAL (XMITQ.FINANCE) USAGE(XMITQ) CLCHNAME(FINANCE.*) REPLACE
```

Figure 16. Changes to isolate the sales queue in a new cluster and separate the gateway cluster transmission queues

4. Remove the SALES queue from the SALES cluster.

Run the **MQSC** command in Figure 17:

```
ALTER QLOCAL(SALES) CLUSTER('Q.SALES') CLUSNL(' ')
```

Figure 17. Remove the sales queue on queue manager SALESRV from the sales cluster

5. Switch the channels to the new transmission queues.

The requirement is to stop and start all the channels that the GATE queue manager is using. To do this with the least number of commands, stop and start the queue manager

```
endmqm -i GATE
strmqm GATE
```

What to do next

1. Rerun the sample request program to verify the new configuration works; see step 2 on page 46
2. Monitor the messages flowing through all the cluster transmission queues on the GATE queue manager:
 - a. Alter the definition of each of the cluster transmission queues to turn queue monitoring on.

```
ALTER QLOCAL(SYSTEM.CLUSTER.TRANSMIT. name) STATQ(ON)
```
 - b. Check queue manager statistics monitoring is OFF, to minimize output, and set the monitoring interval to a lower value to perform multiple tests conveniently.

```
ALTER QMGR STATINT(60) STATCHL(OFF) STATQ(OFF) STATMQI(OFF) STATACLS(OFF)
```
 - c. Restart the GATE queue manager.

- d. Run the sample request program a few times to verify that an equal number of messages are flowing through SYSTEM.CLUSTER.TRANSMIT.Q.SALES.SALESRV and SYSTEM.CLUSTER.TRANSMIT.QUEUE. Requests flow through SYSTEM.CLUSTER.TRANSMIT.Q.SALES.SALESRV and replies through SYSTEM.CLUSTER.TRANSMIT.QUEUE.

```
amqsmon -m GATE -t statistics
```

- e. The results over a couple of intervals are as follows:

```
C:\Documents and Settings\Admin>amqsmon -m GATE -t statistics
MonitoringType: QueueStatistics
QueueManager: 'GATE'
IntervalStartDate: '2012-02-27'
IntervalStartTime: '14.59.20'
IntervalEndDate: '2012-02-27'
IntervalEndTime: '15.00.20'
CommandLevel: 700
ObjectCount: 2
QueueStatistics: 0
QueueName: 'SYSTEM.CLUSTER.TRANSMIT.QUEUE'
CreateDate: '2012-02-24'
CreateTime: '15.58.15'
...
Put1Count: [0, 0]
Put1FailCount: 0
PutBytes: [435, 0]
GetCount: [1, 0]
GetBytes: [435, 0]
...
QueueStatistics: 1
QueueName: 'SYSTEM.CLUSTER.TRANSMIT.Q.SALES.SAVESRV'
CreateDate: '2012-02-24'
CreateTime: '16.37.43'
...
PutCount: [1, 0]
PutFailCount: 0
Put1Count: [0, 0]
Put1FailCount: 0
PutBytes: [435, 0]
GetCount: [1, 0]
GetBytes: [435, 0]
...
MonitoringType: QueueStatistics
QueueManager: 'GATE'
IntervalStartDate: '2012-02-27'
IntervalStartTime: '15.00.20'
IntervalEndDate: '2012-02-27'
IntervalEndTime: '15.01.20'
CommandLevel: 700
ObjectCount: 2
QueueStatistics: 0
QueueName: 'SYSTEM.CLUSTER.TRANSMIT.QUEUE'
CreateDate: '2012-02-24'
CreateTime: '15.58.15'
...
PutCount: [2, 0]
PutFailCount: 0
Put1Count: [0, 0]
```



```

Put1FailCount: 0
PutBytes: [863, 0]
GetCount: [2, 0]
GetBytes: [863, 0]
...
QueueStatistics: 1
QueueName: 'SYSTEM.CLUSTER.TRANSMIT.Q.SALES.SAVESRV'
CreateDate: '2012-02-24'
CreateTime: '16.37.43'
...
PutCount: [2, 0]
PutFailCount: 0
Put1Count: [0, 0]
Put1FailCount: 0
PutBytes: [863, 0]
GetCount: [2, 0]
GetBytes: [863, 0]
...
2 Records Processed.

```

One request and reply message were sent in the first interval and two in the second. You can infer that the request messages were placed on SYSTEM.CLUSTER.TRANSMIT.Q.SALES.SAVESRV, and the reply messages on SYSTEM.CLUSTER.TRANSMIT.QUEUE.

Clustering: Switching cluster transmission queues:

Plan how the changes to the cluster transmission queues of an existing production queue manager are going to be brought into effect.

Before you begin

If you reduce the number of messages the switching process has to transfer to the new transmission queue, switching completes more quickly. Read How the process to switch cluster-sender channel to a different transmission queue works for the reasons for trying to empty the transmission queue before proceeding any further.

About this task

You have a choice of two ways of making the changes to cluster transmission queues take effect.

1. Let the queue manager make the changes automatically. This is the default. The queue manager switches cluster-sender channels with pending transmission queue changes when a cluster-sender channel next starts.
2. Make the changes manually. You can make the changes to a cluster-sender channel when it is stopped. You can switch it from one cluster transmission queue to another before the cluster-sender channel starts.

What factors do you take into account when deciding which of the two options to choose, and how do you manage the switch?

Procedure

- Option 1: Let the queue manager make the changes automatically; see “Switching active cluster-sender channels to another set of cluster-transmission queues” on page 51.

Choose this option if you want the queue manager to make the switch for you.

An alternative way to describe this option is to say the queue manager switches a cluster-sender channel without you forcing the channel to stop. You do have the option of forcing the channel to stop,

and then starting the channel, to make the switch happen sooner. The switch starts when the channel starts, and runs while the channel is running, which is different to option 2. In option 2, the switch takes place when the channel is stopped.

If you choose this option by letting the switch happen automatically, the switching process starts when a cluster-sender channel starts. If the channel is not stopped, it starts after it becomes inactive, if there is a message to process. If the channel is stopped, start it with the `START CHANNEL` command.

The switch process completes as soon as there are no messages left for the cluster-sender channel on the transmission queue the channel was serving. As soon as that is the case, newly arrived messages for the cluster-sender channel are stored directly on the new transmission queue. Until then, messages are stored on the old transmission queue, and the switching process transfers messages from the old transmission queue to the new transmission queue. The cluster-sender channel forwards messages from the new cluster transmission queue during the whole switching process.

When the switch process completes depends on the state of the system. If you are making the changes in a maintenance window, assess beforehand whether the switching process will complete in time. Whether it will complete in time depends on whether the number of messages that are awaiting transfer from the old transmission queue reaches zero.

The advantage of the first method is it is automatic. A disadvantage is that if the time to make the configuration changes is limited to a maintenance window, you must be confident that you can control the system to complete the switch process inside the maintenance window. If you cannot be sure, option 2 might be a better choice.

- Option 2: Make the changes manually; see “Switching a stopped cluster-sender channel to another cluster transmission queue” on page 52.

Choose this option if you want to control the entire switching process manually, or if you want to switch a stopped or inactive channel. It is a good choice, if you are switching a few cluster-sender channels, and you want to do the switch during a maintenance window.

An alternative description of this option is to say that you switch the cluster-sender channel, while the cluster-sender channel is stopped.

If you choose this option you have complete control over when the switch takes place.

You can be certain about completing the switching process in a fixed amount of time, within a maintenance window. The time the switch takes depends on how many messages have to be transferred from one transmission queue to the other. If messages keep arriving, it might take a time for the process to transfer all the messages.

You have the option of switching the channel without transferring messages from the old transmission queue. The switch is “instant”.

When you restart the cluster-sender channel, it starts processing messages on the transmission queue you newly assigned to it.

The advantage of the second method is you have control over the switching process. The disadvantage is that you must identify the cluster-sender channels to be switched, run the necessary commands, and resolve any in-doubt channels that might be preventing the cluster-sender channel stopping.

Related tasks:

“Switching active cluster-sender channels to another set of cluster-transmission queues”

This task gives you three options for switching active cluster-sender channels. One option is to let the queue manager make the switch automatically, which does not affect running applications. The other options are to stop and start channels manually, or to restart the queue manager.

“Switching a stopped cluster-sender channel to another cluster transmission queue” on page 52

Related information:

How the process to switch cluster-sender channel to a different transmission queue works

Switching active cluster-sender channels to another set of cluster-transmission queues:

This task gives you three options for switching active cluster-sender channels. One option is to let the queue manager make the switch automatically, which does not affect running applications. The other options are to stop and start channels manually, or to restart the queue manager.

Before you begin

Change the cluster transmission queue configuration. You can change the **DEFCLXQ** queue manager attribute, or add or modify the **CLCHNAME** attribute of transmission queues.

If you reduce the number of messages the switching process has to transfer to the new transmission queue, switching completes more quickly. Read How the process to switch cluster-sender channel to a different transmission queue works for the reasons for trying to empty the transmission queue before proceeding any further.

About this task

Use the steps in the task as a basis for working out your own plan for making cluster-transmission queue configuration changes.

Procedure

1. Optional: Record the current channel status

Make a record of the status of current and saved channels that are serving cluster transmission queues. The following commands display the status associated with system cluster transmission queues. Add your own commands to display the status associated with cluster-transmission queues that you have defined. Use a convention, such as `XMITQ.ChannelName`, to name cluster transmission queues that you define to make it easy to display the channel status for those transmission queues.

```
DISPLAY CHSTATUS(*) WHERE(XMITQ LK 'SYSTEM.CLUSTER.TRANSMIT.*')
DISPLAY CHSTATUS(*) SAVED WHERE(XMITQ LK 'SYSTEM.CLUSTER.TRANSMIT.*')
```

2. Switch transmission queues.

- Do nothing. The queue manager switches cluster-sender channels when they restart after being stopped or inactive.

Choose this option if you have no rules or concerns about altering a queue manager configuration. Running applications are not affected by the changes.

- Restart the queue manager. All cluster-sender channels are stopped and restarted automatically on demand.

Choose this option to initiate all the changes immediately. Running applications are interrupted by the queue manager as it shuts down and restarts.

- Stop individual cluster-sender channels and restart them.

Choose this option to change a few channels immediately. Running applications experience a short delay in message transfer between your stopping and starting the message channel again. The cluster-sender channel remains running, except during the time you stopped it. During the switch

process messages are delivered to the old transmission queue, transferred to the new transmission queue by the switching process, and forwarded from the new transmission queue by the cluster-sender channel.

3. Optional: Monitor the channels as they switch

Display the channel status and the transmission queue depth during the switch. The following example display the status for system cluster transmission queues.

```
DISPLAY CHSTATUS(*) WHERE(XMITQ LK 'SYSTEM.CLUSTER.TRANSMIT.*')
DISPLAY CHSTATUS(*) SAVED WHERE(XMITQ LK 'SYSTEM.CLUSTER.TRANSMIT.*')
DISPLAY QUEUE('SYSTEM.CLUSTER.TRANSMIT.*') CURDEPTH
```

4. Optional: Monitor the messages “ AMQ7341 The transmission queue for channel *ChannelName* switched from queue *QueueName* to *QueueName* ” that are written to the queue manager error log.

Switching a stopped cluster-sender channel to another cluster transmission queue:

Before you begin

You might make some configuration changes, and now want to make them effective without starting the cluster-sender channels that are affected. Alternatively, you make the configuration changes you require as one of the steps in the task.

If you reduce the number of messages the switching process has to transfer to the new transmission queue, switching completes more quickly. Read How the process to switch cluster-sender channel to a different transmission queue works for the reasons for trying to empty the transmission queue before proceeding any further.

About this task

This task switches the transmission queues served by stopped or inactive cluster-sender channels. You might do this task because a cluster-sender channel is stopped, and you want to switch its transmission queue immediately. For example, for some reason a cluster-sender channel is not starting, or has some other configuration problem. To resolve the problem, you decide to create a cluster-sender channel, and associate the transmission queue for the old cluster-sender channel with the new cluster-sender channel you defined.

A more likely scenario is you want to control when reconfiguration of cluster transmission queues is performed. To fully control the reconfiguration, you stop the channels, change the configuration, and then switch the transmission queues.

Procedure

1. Stop the channels that you intend to switch

- a. Stop any running or inactive channels that you intend to switch. Stopping an inactive cluster-sender channel prevents it starting while you are making configuration changes.

```
STOP CHANNEL(ChannelName) MODE(QUIESCSE) STATUS(STOPPED)
```


2. Optional: Make the configuration changes.

For example, see “Clustering: Example configuration of multiple cluster transmission queues” on page 42.

3. Switch the cluster-sender channels to the new cluster transmission queues.

 On distributed platforms, issue the following command:

```
runswchl -m QmgrName -c ChannelName
```

 On z/OS, use the SWITCH function of the CSQUTIL command to switch the messages or monitor what is happening. Use the following command.

```
SWITCH CHANNEL(channel_name) MOVEMSGS(YES)
```

For more information, see SWITCH function.

The **runswchl**, or CSQUTIL SWITCH, command transfers any messages on the old transmission queue to the new transmission queue. When the number of messages on the old transmission queue for this channel reaches zero, the switch is completed. The command is synchronous. The command writes progress messages to the window during the switching process.

During the transfer phase existing and new messages destined for the cluster-sender channel are transferred in order to the new transmission queue.

Because the cluster-sender channel is stopped, the messages build up on the new transmission queue. Contrast the stopped cluster-sender channel, to step 2 on page 51 in “Switching active cluster-sender channels to another set of cluster-transmission queues” on page 51. In that step, the cluster-sender channel is running, so messages do not necessarily build up on the new transmission queue.

4. Optional: Monitor the channels as they switch

In a different command window, display the transmission queue depth during the switch. The following example display the status for system cluster transmission queues.

```
DISPLAY QUEUE('SYSTEM.CLUSTER.TRANSMIT.*') CURDEPTH
```

5. Optional: Monitor the messages “ AMQ7341 The transmission queue for channel *ChannelName* switched from queue *QueueName* to *QueueName* ” that are written to the queue manager error log.

6. Restart the cluster-sender channels that you stopped.

The channels do not start automatically, as you stopped them, placing them into STOPPED status.

```
START CHANNEL(ChannelName)
```

Related information:

runswchl

RESOLVE CHANNEL

STOP CHANNEL

Clustering: Migration and modification best practices:

This topic provides guidance for planning and administering IBM MQ clusters. This information is a guide based on testing and feedback from customers.

1. “Moving objects in a cluster” (Best practices for moving objects around inside a cluster, without installing any fix packs or new versions of IBM MQ).
2. “Upgrades and maintenance installations” on page 55 (Best practices for keeping a working cluster architecture up and running, while applying maintenance or upgrades and testing the new architecture).

Moving objects in a cluster

Applications and their queues

When you must move a queue instance hosted on one queue manager to be hosted on another, you can work with the workload balancing parameters to ensure a smooth transition.

Create an instance of the queue where it is to be newly hosted, but use cluster workload balancing settings to continue sending messages to the original instance until your application is ready to switch. This is achieved with the following steps:

1. Set the **CLWLRANK** property of the existing queue to a high value, for example five.
2. Create the new instance of the queue and set its **CLWLRANK** property to zero.
3. Complete any further configuration of the new system, for example deploy and start consuming applications against the new instance of the queue.
4. Set the **CLWLRANK** property of the new queue instance to be higher than the original instance, for example nine.

5. Allow the original queue instance to process any queued messages in the system and then delete the queue.

Moving entire queue managers

If the queue manager is staying on the same host, but the IP address is changing, then the process is as follows:

- DNS, when used correctly, can help simplify the process. For information about using DNS by setting the Connection name (CONNNAME) channel attribute, see ALTER CHANNEL.
- If moving a full repository, ensure that you have at least one other full repository which is running smoothly (no problems with channel status for example) before making changes.
- Suspend the queue manager using the SUSPEND QMGR command to avoid traffic buildup.
- Modify the IP address of the computer. If your CLUSRCVR channel definition uses an IP address in the CONNNAME field, modify this IP address entry. The DNS cache might need to be flushed through to ensure that updates are available everywhere.
- When the queue manager reconnects to the full repositories, channel auto-definitions automatically resolve themselves.
- If the queue manager hosted a full repository and the IP address changes, it is important to ensure that partials are switched over as soon as possible to point any manually defined CLUSSDR channels to the new location. Until this switch is carried out, these queue managers might be able to only contact the remaining (unchanged) full repository, and warning messages might be seen regarding the incorrect channel definition.
- Resume the queue manager using the RESUME QMGR command.

If the queue manager must be moved to a new host, it is possible to copy the queue manager data and restore from a backup. This process is not recommended however, unless there are no other options; it might be better to create a queue manager on a new machine and replicate queues and applications as described in the previous section. This situation gives a smooth rollover/rollback mechanism.

If you are determined to move a complete queue manager using backup, follow these best practices:

- Treat the whole process as a queue manager restore from backup, applying any processes you would usually use for system recovery as appropriate for your operating system environment.
- Use the **REFRESH CLUSTER** command after migration to discard all locally held cluster information (including any auto-defined channels that are in doubt), and force it to be rebuilt.

Note: For large clusters, using the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.


When creating a queue manager and replicating the setup from an existing queue manager in the cluster (as described previously in this topic), never treat the two different queue managers as actually being the same. In particular, do not give a new queue manager the same queue manager name and IP address. Attempting to 'drop in' a replacement queue manager is a frequent cause of problems in IBM MQ clusters. The cache expects to receive updates including the **QMID** attribute, and state can be corrupted.

If two different queue managers are accidentally created with the same name, it is recommended to use the **RESET CLUSTER QMID** command to eject the incorrect entry from the cluster.

Upgrades and maintenance installations

Avoid the "big bang scenario" (for example, stopping all cluster and queue manager activity, applying all upgrades and maintenance to all queue managers, then starting everything at the same time): Clusters are designed to still work with multiple versions of queue manager coexisting, so a well-planned, phased maintenance approach is recommended.

Have a backup plan:

-  On z/OS, have you applied backwards migration PTFs?
- Have you taken backups?
- Avoid using new cluster functionality immediately: Wait until you are sure that all the queue managers are upgraded to the new level, and are certain that you are not going to roll any of them back. Using new cluster function in a cluster where some queue managers are still at an earlier level can lead to undefined behavior. For example, in the move to IBM WebSphere® MQ Version 7.1 from IBM WebSphere MQ Version 6.0, if a queue manager defines a cluster topic, IBM WebSphere MQ Version 6.0 queue managers will not understand the definition or be able to publish on this topic.

Migrate full repositories first. Although they can forward information that they do not understand, they cannot persist it, so it is not the recommended approach unless absolutely necessary. For more information, see Queue manager cluster migration.

Clustering: Using REFRESH CLUSTER best practices:

You use the **REFRESH CLUSTER** command to discard all locally held information about a cluster and rebuild that information from the full repositories in the cluster. You should not need to use this command, except in exceptional circumstances. If you do need to use it, there are special considerations about how you use it. This information is a guide based on testing and feedback from customers.

Only run REFRESH CLUSTER if you really need to do so

The IBM MQ cluster technology ensures that any change to the cluster configuration, such as a change to a clustered queue, automatically becomes known to any member of the cluster that needs to know the information. There is no need for further administrative steps to be taken to achieve this propagation of information.

If such information does not reach the queue managers in the cluster where it is required, for example a clustered queue is not known by another queue manager in the cluster when an application attempts to open it for the first time, it implies a problem in the cluster infrastructure. For example, it is possible that a channel cannot be started between a queue manager and a full repository queue manager. Therefore, any situation where inconsistencies are observed must be investigated. If possible, resolve the situation without using the **REFRESH CLUSTER** command.

In rare circumstances that are documented elsewhere in this product documentation, or when requested by IBM support, you can use the **REFRESH CLUSTER** command to discard all locally held information about a cluster and rebuild that information from the full repositories in the cluster.

Refreshing in a large cluster can affect performance and availability of the cluster

Use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, for example by creating a sudden increase in work for the full repositories as they process the repropagation of queue manager cluster resources. If you are refreshing in a large cluster (that is, many hundreds of queue managers) you should avoid use of the command in day-to-day work if possible and use alternative methods to correct specific inconsistencies. For example, if a cluster queue is not being correctly propagated across the cluster, an initial investigation technique of updating the clustered queue

definition, such as altering its description, repropagates the queue configuration across the cluster. This process can help to identify the problem and potentially resolve a temporary inconsistency.

If alternative methods cannot be used, and you have to run **REFRESH CLUSTER** in a large cluster, you should do so at off-peak times or during a maintenance window to avoid impact on user workloads. You should also avoid refreshing a large cluster in a single batch, and instead stagger the activity as explained in “Avoid performance and availability issues when cluster objects send automatic updates.”

Avoid performance and availability issues when cluster objects send automatic updates

After a new cluster object is defined on a queue manager, an update for this object is generated every 27 days from the time of definition, and sent to every full repository in the cluster and onwards to any other interested queue managers. When you issue the **REFRESH CLUSTER** command to a queue manager, you reset the clock for this automatic update on all objects defined locally in the specified cluster.

If you refresh a large cluster (that is, many hundreds of queue managers) in a single batch, or in other circumstances such as recreating a system from configuration backup, after 27 days all of those queue managers will re-advertise all of their object definitions to the full repositories at the same time. This could again cause the system to run significantly slower, or even become unavailable, until all the updates have completed. Therefore, when you have to refresh or recreate multiple queue managers in a large cluster, you should stagger the activity over several hours, or several days, so that subsequent automatic updates do not regularly impact system performance.

The system cluster history queue

When a **REFRESH CLUSTER** is performed, the queue manager takes a snapshot of the cluster state before the refresh and stores it on the `SYSTEM.CLUSTER.HISTORY.QUEUE (SCHQ)` if it is defined on the queue manager. This snapshot is for IBM service purposes only, in case of later problems with the system.

The SCHQ is defined by default on distributed queue managers on startup. For z/OS migration, the SCHQ must be manually defined.

Messages on the SCHQ expire after three months.

Related information:

REFRESH CLUSTER considerations for publish/subscribe clusters

Clustering: Availability, multi instance, and disaster recovery:

This topic provides guidance for planning and administering IBM MQ clusters. This information is a guide based on testing and feedback from customers.

IBM MQ Clustering itself is not a High Availability solution, but in some circumstances it can be used to improve the availability of services using IBM MQ, for example by having multiple instances of a queue on different queue managers. This section gives guidance on ensuring that the IBM MQ infrastructure is as highly available as possible so that it can be used in such an architecture.

Availability of cluster resources

The reason for the usual recommendation to maintain two full repositories is that the loss of one is not critical to the smooth running of the cluster. Even if both become unavailable, there is a 60 day grace period for existing knowledge held by partial repositories, although new or not previously accessed resources (queues for example) are not available in this event.

Using clusters to improve application availability

A cluster can help in designing highly available applications (for example a request/response type server application), by using multiple instances of the queue and application. If needed, priority attributes can give preference to the 'live' application unless a queue manager or channel

for example become unavailable. This is powerful for switching over quickly to continue processing new messages when a problem occurs.

However, messages which were delivered to a particular queue manager in a cluster are held only on that queue instance, and are not available for processing until that queue manager is recovered. For this reason, for true data high availability you might want to consider other technologies such as multi-instance queue managers.


Multi-instance queue managers

Software High Availability (multi instance) is the best built-in offering for keeping your existing messages available. See Using IBM MQ with high availability configurations, Create a multi-instance queue manager, and the following section for more information. Any queue manager in a cluster may be made highly available using this technique, as long as all queue managers in the cluster are running at least IBM WebSphere MQ Version 7.0.1. If any queue managers in the cluster are at previous levels, they might lose connectivity with the multi-instance queue managers if they fail over to a secondary IP.

As discussed previously in this topic, as long as two full repositories are configured, they are almost by their nature highly available. If you need to, IBM MQ software High Availability / multi-instance queue managers can be used for full repositories. There is no strong reason to use these methods, and in fact for temporary outages these methods might cause additional performance cost during the failover. Using software HA instead of running two full repositories is discouraged because in the event of a single channel outage, for example, it would not necessarily fail over, but might leave partial repositories unable to query for cluster resources.

Disaster recovery

Disaster recovery, for example recovering from when the disks storing a queue manager's data becomes corrupt, is difficult to do well; IBM MQ can help, but it cannot do it automatically. The only 'true' disaster recovery option in IBM MQ (excluding any operating system or other underlying replication technologies) is restoration from a backup. There are some cluster specific points to consider in these situations:

- Take care when testing disaster recovery scenarios. For example, if testing the operation of backup queue managers, be careful when bringing them online in the same network as it is possible to accidentally join the live cluster and start 'stealing' messages by hosting the same named queues as in the live cluster queue managers.
- Disaster recovery testing must not interfere with a running live cluster. Techniques to avoid interference include:
 - Complete network separation or separation at the firewall level.
 -  Not starting channel initiation or the z/OS **chinit** address space.
 - Not issuing live SSL certificate to the disaster recovery system until, or unless, an actual disaster recovery scenario occurs.
- When restoring a backup of a queue manager in the cluster it is possible that the backup is out of sync with the rest of the cluster. The **REFRESH CLUSTER** command can resolve updates and synchronize with the cluster but the **REFRESH CLUSTER** command must be used as a last resort. See "Clustering: Using REFRESH CLUSTER best practices" on page 55. Review any in-house process documentation and IBM MQ documentation to see whether a simple step was missed before resorting to using the command.
- As for any recovery, applications must deal with replay and loss of data. It must be decided whether to clear the queues down to a known state, or if there is enough information elsewhere to manage replays.

Planning your distributed publish/subscribe network

You can create a network of queue managers where subscriptions created on one queue manager will receive matching messages published by an application connected to another queue manager in the network. To choose a suitable topology, you need to consider your requirements for manual control, network size, frequency of change, availability and scalability.

Before you begin

This task assumes that you understand what distributed publish/subscribe networks are, and how they work. For a technical overview, see *Distributed publish/subscribe networks*.

About this task

There are three basic topologies for a publish/subscribe network:

- Direct routed cluster
- Topic host routed cluster
- Hierarchy

For the first two topologies, the starting point is an IBM MQ cluster configuration. The third topology can be created with or without a cluster. See “Planning your distributed queues and clusters” on page 3 for information about planning the underlying queue manager network.

A *Direct routed cluster* is the simplest topology to configure when a cluster is already present. Any topic that you define on any queue manager is automatically made available on every queue manager in the cluster, and publications are routed directly from any queue manager where a publishing application connects, to each of the queue managers where matching subscriptions exist. This simplicity of configuration relies on IBM MQ maintaining a high level of sharing of information and connectivity between every queue manager in the cluster. For small and simple networks (that is, a small number of queue managers, and a fairly static set of publishers and subscribers) this is acceptable. However, when used in larger or more dynamic environments the overhead might be prohibitive. See “Direct routing in publish/subscribe clusters” on page 63.

A *Topic host routed cluster* gives the same benefit as a direct routed cluster, by making any topic that you define on any queue manager in the cluster automatically available on every queue manager in the cluster. However, topic host routed clusters require you to carefully choose the queue managers that host each topic, because all information and publications for that topic pass through those topic host queue managers. This means that the system does not have to maintain channels and information flows between all queue managers. However it also means that publications might no longer be sent directly to subscribers, but might be routed through a topic host queue manager. For these reasons additional load might be put on the system, especially on the queue managers hosting the topics, so careful planning of the topology is required. This topology is particularly effective for networks that contain many queue managers, or that host a dynamic set of publishers and subscribers (that is, publishers or subscribers that are frequently added or removed). Additional topic hosts can be defined to improve availability of routes and to horizontally scale publication workload. See “Topic host routing in publish/subscribe clusters” on page 68.

A *Hierarchy* requires the most manual configuration to set up, and is the hardest topology to modify. You must manually configure the relationships between each queue manager in the hierarchy and its direct relations. After relationships are configured, publications will (as for the previous two topologies) be routed to subscriptions on other queue managers in the hierarchy. Publications are routed using the hierarchy relationships. This allows very specific topologies to be configured to suit different requirements, but it can also result in publications requiring many “hops” through intermediate queue managers to reach the subscriptions. There is always only one route through a hierarchy for a publication, so availability of every queue manager is critical. Hierarchies are typically only preferable

where a single cluster cannot be configured; for example when spanning multiple organizations. See “Routing in publish/subscribe hierarchies” on page 93.

Where necessary, the above three topologies can be combined to solve specific topographical requirements. For an example, see Combining the topic spaces of multiple clusters.

To choose a suitable topology for your distributed publish/subscribe network, you need to consider the following broad questions:

- How big will your network be?
- How much manual control do you need over its configuration?
- How dynamic will the system be, both in terms of topics and subscriptions, and in terms of queue managers?
- What are your availability and scalability requirements?
- Can all queue managers connect directly to each other?

Procedure

- Estimate how big your network needs to be.
 1. Estimate how many topics you need.
 2. Estimate how many publishers and subscribers you expect to have.
 3. Estimate how many queue managers will be involved in publish/subscribe activities.
See also “Publish/subscribe clustering: Best practices” on page 78, especially the following sections:
 - How to size your system
 - Reasons to limit the number of cluster queue managers involved in publish/subscribe activity
 - How to decide which topics to cluster

If your network will have many queue managers, and handle many publishers and subscribers, you probably need to use a topic host routed cluster or a hierarchy. Direct routed clusters require almost no manual configuration, and can be a good solution for small or static networks.

- Consider how much manual control you need over which queue manager hosts each topic, publisher or subscriber.
 1. Consider whether some of your queue managers are less capable than others.
 2. Consider whether the communication links to some of your queue managers are more fragile than to others.
 3. Identify cases where you expect a topic to have many publications and few subscribers.
 4. Identify cases where you expect a topic to have many subscribers and few publications.

In all topologies, publications are delivered to subscriptions on other queue managers. In a direct routed cluster those publications take the shortest path to the subscriptions. In a topic host routed cluster or a hierarchy, you control the route that publications take. If your queue managers differ in their capability, or have differing levels of availability and connectivity, you will probably want to assign specific workloads to specific queue managers. You can do this using either a topic host routed cluster or a hierarchy.

In all topologies, co-locating the publishing applications on the same queue manager as the subscriptions whenever possible minimizes overheads and maximizes performance. For topic host routed clusters, consider putting publishers or subscribers on the queue managers that host the topic. This removes any extra “hops” between queue managers to pass a publication to a subscriber. This approach is particularly effective in cases where a topic has many publishers and few subscribers, or many subscribers and few publishers. See, for example, Topic host routing using centralized publishers or subscribers.

- See also “Publish/subscribe clustering: Best practices” on page 78, especially the following sections:
- How to decide which topics to cluster
 - Publisher and subscription location

- Consider how dynamic the network activity will be.

1. Estimate how frequently subscribers will be added and removed on different topics.

Whenever a subscription is added or removed from a queue manager, and it is the first or last subscription for that specific topic string, that information is communicated to other queue managers in the topology. In a direct routed cluster and a hierarchy, this subscription information is propagated to every queue manager in the topology whether or not they have publishers on the topic. If the topology consists of many queue managers, this might be a significant performance overhead. In a topic host routed cluster, this information is only propagated to those queue managers that host a clustered topic that maps to the subscription's topic string.

See also the Subscription change and dynamic topic strings section of “Publish/subscribe clustering: Best practices” on page 78.

Note: In very dynamic systems, where the set of many unique topic strings is rapidly and constantly being changed, it might be best to switch the model to a “publish everywhere” mode. See Subscription performance in publish/subscribe networks.

2. Consider how dynamic the queue managers are in the topology.

A hierarchy requires each change in queue manager in the topology to be manually inserted or removed from the hierarchy, with care taken when changing queue managers at higher levels in the hierarchy. Queue managers in a hierarchy typically also use manually configured channel connections. You must maintain these connections, adding and removing channels as queue managers are added and removed from the hierarchy.

In a publish/subscribe cluster, queue managers are automatically connected to any other queue manager that is required when they first join the cluster, and automatically become aware of topics and subscriptions.

- Consider your route availability and publication traffic scalability requirements.

1. Decide whether you need to always have an available route from a publishing queue manager to a subscribing queue manager, even when a queue manager is unavailable.

2. Consider how scalable you need the network to be. Decide whether the level of publication traffic is too high to be routed through a single queue manager or channel, and whether that level of publication traffic must be handled by a single topic branch or can be spread across multiple topic branches.

3. Consider whether you need to maintain message ordering.

Because a direct routed cluster sends messages directly from publishing queue managers to subscribing queue managers, you do not need to consider the availability of intermediate queue managers along the route. Similarly, scaling to the intermediate queue managers is not a consideration. However, as previously mentioned, the overhead of automatically maintaining channels and information flows between all queue managers in the cluster can significantly affect performance, especially in a large or dynamic environment.

A topic host routed cluster can be tuned for individual topics. You can ensure that each branch of the topic tree that has a considerable publication workload is defined on a different queue manager, and that each queue manager is sufficiently performant and available for the expected workload for that branch of the topic tree. You can also improve availability and horizontal scaling further by defining each topic on multiple queue managers. This allows the system to route around unavailable topic host queue managers, and to workload balance publication traffic across them. However, when you define a given topic on multiple queue managers, you also introduce the following constraints:

- You lose message ordering across publications.
- You cannot use retained publications. See “Design considerations for retained publications in publish/subscribe clusters” on page 91.

You cannot configure high availability or scalability of routing in a hierarchy through multiple routes.

See also the Publication traffic section of “Publish/subscribe clustering: Best practices” on page 78.

- Based on these calculations, use the links provided to help you decide whether to use a topic host routed cluster, a direct routed cluster, a hierarchy, or a mixture of these topologies.

What to do next

You are now ready to configure your distributed publish/subscribe network.

Related information:

- Configuring a queue manager cluster
- Configuring distributed queuing
- Configuring a publish/subscribe cluster
- Connecting a queue manager to a publish/subscribe hierarchy

Designing publish/subscribe clusters

There are two basic publish/subscribe cluster topologies: *direct routing* and *topic host routing*. Each has different benefits. When you design your publish/subscribe cluster, choose the topology that best fits your expected network requirements.

For an overview of the two publish/subscribe cluster topologies, see *Publish/subscribe clusters*. To help you evaluate your network requirements, see “Planning your distributed publish/subscribe network” on page 58 and “Publish/subscribe clustering: Best practices” on page 78.

In general, both cluster topologies provide the following benefits:

- Simple configuration on top of a point-to-point cluster topology.
- Automatic handling of queue managers joining and leaving the cluster.
- Ease of scaling for additional subscriptions and publishers, by adding extra queue managers and distributing the additional subscriptions and publishers across them.

However, the two topologies have different benefits as the requirements become more specific.

Direct routed publish/subscribe clusters

With direct routing, any queue manager in the cluster sends publications from connected applications direct to any other queue manager in the cluster with a matching subscription.

A direct routed publish/subscribe cluster provides the following benefits:

- Messages destined for a subscription on a specific queue manager in the same cluster are transported directly to that queue manager and do not need to pass through an intermediate queue manager. This can improve performance in comparison with a topic host routed topology, or a hierarchical topology.
- Because all queue managers are directly connected to each other, there is no single point of failure in the routing infrastructure of this topology. If one queue manager is not available, subscriptions on other queue managers in the cluster are still able to receive messages from publishers on available queue managers.
- It is very simple to configure, especially on an existing cluster.

Things to consider when using a direct routed publish/subscribe cluster:

- All queue managers in the cluster become aware of all other queue managers in the cluster.
- Queue managers in a cluster that host one or more subscriptions to a clustered topic, automatically create cluster sender channels to all other queue managers in the cluster, even when those queue managers are not publishing messages on any clustered topics.
- The first subscription on a queue manager to a topic string under a clustered topic results in a message being sent to every other queue manager in the cluster. Similarly, the last subscription on a topic string to be deleted also results in a message. The more individual topic strings being used under a clustered topic, and the higher the rate of change of subscriptions, the more inter-queue manager communication occurs.

- Every queue manager in the cluster maintains the knowledge of subscribed topic strings that it is informed of, even when the queue manager is neither publishing nor subscribing to those topics.

For the above reasons, all queue managers in a cluster with a direct routed topic defined will incur an additional overhead. The more queue managers there are in the cluster, the greater the overhead. Likewise the more topic strings subscribed to, and the greater their rate of change, the greater the overhead. This can result in too much load on queue managers running on small systems in a large or dynamic direct routed publish/subscribe cluster. See *Direct routed publish/subscribe performance* for further information.

When you know that a cluster cannot accommodate the overheads of direct routed clustered publish/subscribe, you can instead use topic host routed publish/subscribe. Alternatively, in extreme situations, you can completely disable clustered publish/subscribe functionality by setting the queue manager attribute **PSCLUS** to **DISABLED** on every queue manager in the cluster. See “Inhibiting clustered publish/subscribe” on page 88. This prevents any clustered topic from being created, and therefore ensures that your network does not incur any overheads associated with clustered publish/subscribe.

Topic host routed publish/subscribe clusters

With topic host routing, the queue managers where clustered topics are administratively defined become routers for publications. Publications from non-hosting queue managers in the cluster are routed through the hosting queue manager to any queue manager in the cluster with a matching subscription.

A topic host routed publish/subscribe cluster provides the following extra benefits over a direct routed publish/subscribe cluster:

- Only queue managers on which topic host routed topics are defined are made aware of all other queue managers in the cluster.
- Only the topic host queue managers need to be able to connect to all other queue managers in the cluster, and will typically only connect to those where subscriptions exist. Therefore there are significantly fewer channels running between queue managers.
- Cluster queue managers that host one or more subscriptions to a clustered topic automatically create cluster sender channels only to queue managers that host a cluster topic that maps to the topic string of the subscription.
- The first subscription on a queue manager to a topic string under a clustered topic results in a message being sent to a queue manager in the cluster that hosts the clustered topic. Similarly, the last subscription on a topic string to be deleted also results in a message. The more individual topic strings being used under a clustered topic, and the higher the rate of change of subscriptions, the more inter-queue manager communication occurs, but only between subscription hosts and topic hosts.
- More control over the physical configuration. With direct routing, all queue managers have to participate in the publish/subscribe cluster, increasing their overheads. With topic host routing, only the topic host queue managers are aware of other queue managers and their subscriptions. You explicitly choose the topic host queue managers, therefore you can ensure that those queue managers are running on adequate equipment, and you can use less powerful systems for the other queue managers.

Things to consider when using a topic host routed publish/subscribe cluster:

- An extra “hop” between a publishing queue manager and a subscribing queue manager is introduced when the publisher or the subscriber is not located on a topic hosting queue manager. The latency caused by the extra “hop” can mean that topic host routing is less efficient than direct routing.
- On large clusters, topic host routing eases the significant performance and scaling issues that you can get with direct routing.
- You might choose to define all your topics on a single queue manager, or on a very small number of queue managers. If you do this, make sure the topic host queue managers are hosted on powerful systems with good connectivity.

- You can define the same topic on more than one queue manager. This improves the availability of the topic, and also improves scalability because IBM MQ workload balances publications for a topic across all hosts for that topic. Note, however, that defining the same topic on more than one queue manager loses message order for that topic.
- By hosting different topics on different queue managers, you can improve scalability without losing message order.

Related information:

Publish/subscribe cluster scenario

Configuring a publish/subscribe cluster

Tuning distributed publish/subscribe networks

Distributed publish/subscribe troubleshooting

Direct routing in publish/subscribe clusters:

Publications from any publishing queue manager are routed direct to any other queue manager in the cluster with a matching subscription.

For an introduction to how messages are routed between queue managers in publish/subscribe hierarchies and clusters, see Distributed publish/subscribe networks.

A direct routed publish/subscribe cluster behaves as follows:

- All queue managers automatically know of all other queue managers.
- All queue managers with subscriptions to clustered topics create channels to all other queue managers in the cluster and inform them of their subscriptions.
- Messages published by an application are routed from the queue manager that it is connected to, direct to each queue manager where a matching subscription exists.

The following diagram shows a queue manager cluster that is not currently used for publish/subscribe or point-to-point activities. Note that every queue manager in the cluster connects only to and from the full repository queue managers.




Figure 18. A queue manager cluster

For publications to flow between queue managers in a direct routed cluster, you cluster a branch of the topic tree as described in *Configuring a publish/subscribe cluster*, and specify *direct routing* (the default).

In a direct routed publish/subscribe cluster, you define the topic object on any queue manager in the cluster. When you do this, knowledge of the object, and knowledge of all other queue managers in the cluster, is automatically pushed to all queue managers in the cluster by the full repository queue managers. This happens before any queue manager references the topic:

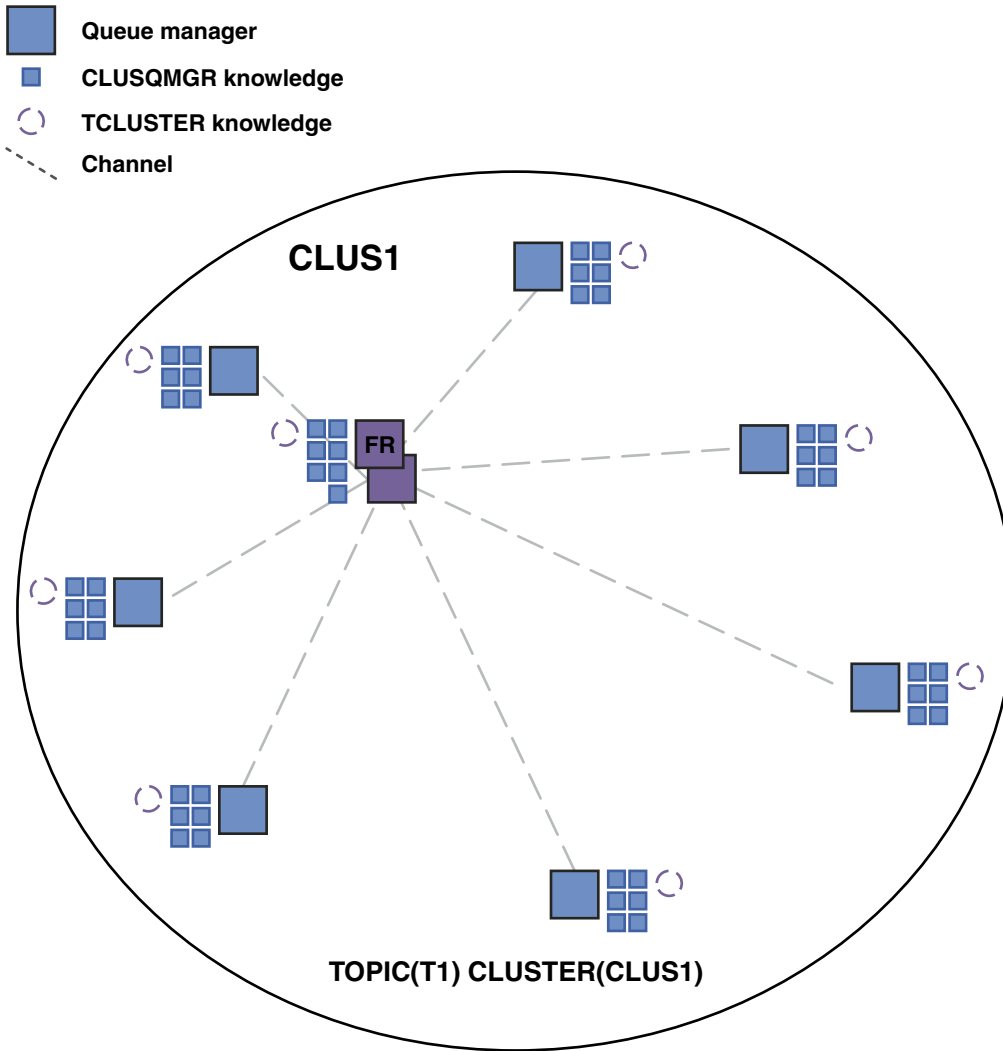


Figure 19. A direct routed publish/subscribe cluster

When a subscription is created, the queue manager that hosts the subscription establishes a channel to every queue manager in the cluster, and sends details of the subscription. This distributed subscription knowledge is represented by a proxy subscription on each queue manager. When a publication is produced on any queue manager in the cluster that matches that proxy subscription's topic string, a cluster channel is established from the publisher queue manager to each queue manager hosting a subscription, and the message is sent to each of them.

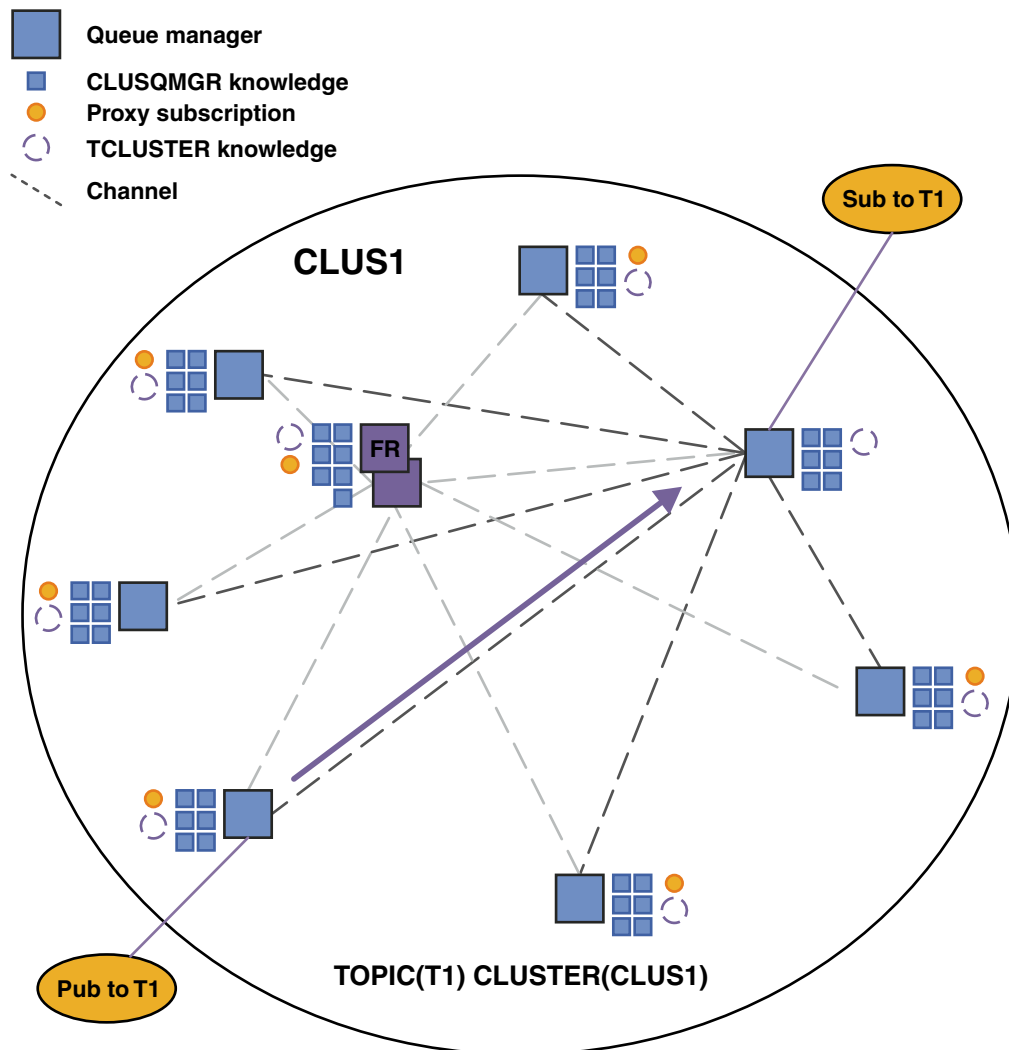


Figure 20. A direct routed publish/subscribe cluster with a publisher and a subscriber to a clustered topic

The direct routing of publications to subscription hosting queue managers simplifies configuration and minimizes the latency in delivering publications to subscriptions.

However, depending on the location of subscriptions and publishers, your cluster can quickly become fully interconnected, with every queue manager having a direct connection to every other queue manager. This might or might not be acceptable in your environment. Similarly, if the set of topic strings being subscribed to is changing frequently, the overhead of propagating that information between all queue managers can also become significant. All queue managers in a direct routed publish/subscribe cluster must be able to cope with these overheads.

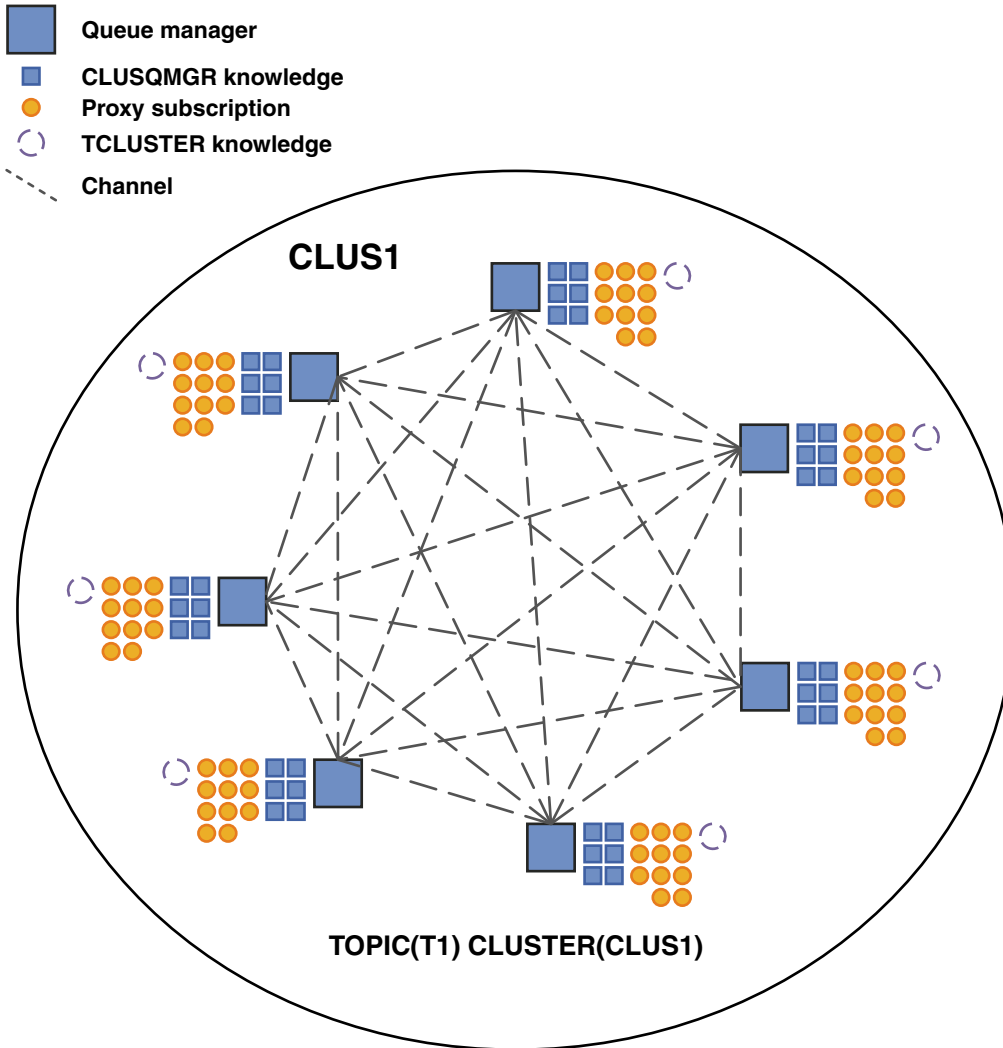


Figure 21. A direct routed publish/subscribe cluster that is fully interconnected

Summary and additional considerations

A direct routed publish/subscribe cluster needs little manual intervention to create or administer, and provides direct routing between publishers and subscribers. For certain configurations it is usually the most appropriate topology, notably clusters with few queue managers, or where high queue manager connectivity is acceptable and subscriptions are changing infrequently. However it also imposes certain constraints upon your system:

- The load on each queue manager is proportional to the total number of queue managers in the cluster. Therefore, in larger clusters, individual queue managers and the system as a whole can experience performance issues.
- By default, all clustered topic strings subscribed to are propagated throughout the cluster, and publications are propagated only to remote queue managers that have a subscription to the associated topic. Therefore rapid changes to the set of subscriptions can become a limiting factor. You can change this default behavior, and instead have all publications propagated to all queue managers, which removes the need for proxy subscriptions. This reduces the subscription knowledge traffic, but is likely to increase the publication traffic and the number of channels each queue manager establishes. See Subscription performance in publish/subscribe networks.

Note: A similar restriction also applies to hierarchies.

- Because of the interconnected nature of publish/subscribe queue managers, it takes time for proxy subscriptions to propagate around all nodes in the network. Remote publications do not necessarily start being subscribed to immediately, so early publications might not be sent following a subscription to a new topic string. You can remove the problems caused by the subscription delay by having all publications propagated to all queue managers, which removes the need for proxy subscriptions. See Subscription performance in publish/subscribe networks.

Note: This restriction also applies to hierarchies.

Before you use direct routing, explore the alternative approaches detailed in “Topic host routing in publish/subscribe clusters,” and “Routing in publish/subscribe hierarchies” on page 93.

Topic host routing in publish/subscribe clusters:

Publications from non-hosting queue managers in the cluster are routed through the hosting queue manager to any queue manager in the cluster with a matching subscription.

For an introduction to how messages are routed between queue managers in publish/subscribe hierarchies and clusters, see Distributed publish/subscribe networks.

To understand the behavior and benefits of topic host routing it is best to first understand “Direct routing in publish/subscribe clusters” on page 63.

A topic host routed publish/subscribe cluster behaves as follows:

- Clustered administered topic objects are manually defined on individual queue managers in the cluster. These are referred to as *topic host queue managers*.
- When a subscription is made on a cluster queue manager, channels are created from the subscription host queue manager to the topic host queue managers, and proxy subscriptions are created only on the queue managers that host the topic.
- When an application publishes information to a topic, the connected queue manager always forwards the publication to one queue manager that hosts the topic, which passes it on to all queue managers in the cluster that have matching subscriptions to the topic.

This process is explained in more detail in the following examples.

Topic host routing using a single topic host

For publications to flow between queue managers in a topic host routed cluster, you cluster a branch of the topic tree as described in Configuring a publish/subscribe cluster, and specify *topic host routing*.

There are a number of reasons to define a topic host routed topic object on multiple queue managers in a cluster. However, for simplicity we start with a single topic host.

The following diagram shows a queue manager cluster that is not currently used for publish/subscribe or point-to-point activities. Note that every queue manager in the cluster connects only to and from the full repository queue managers.

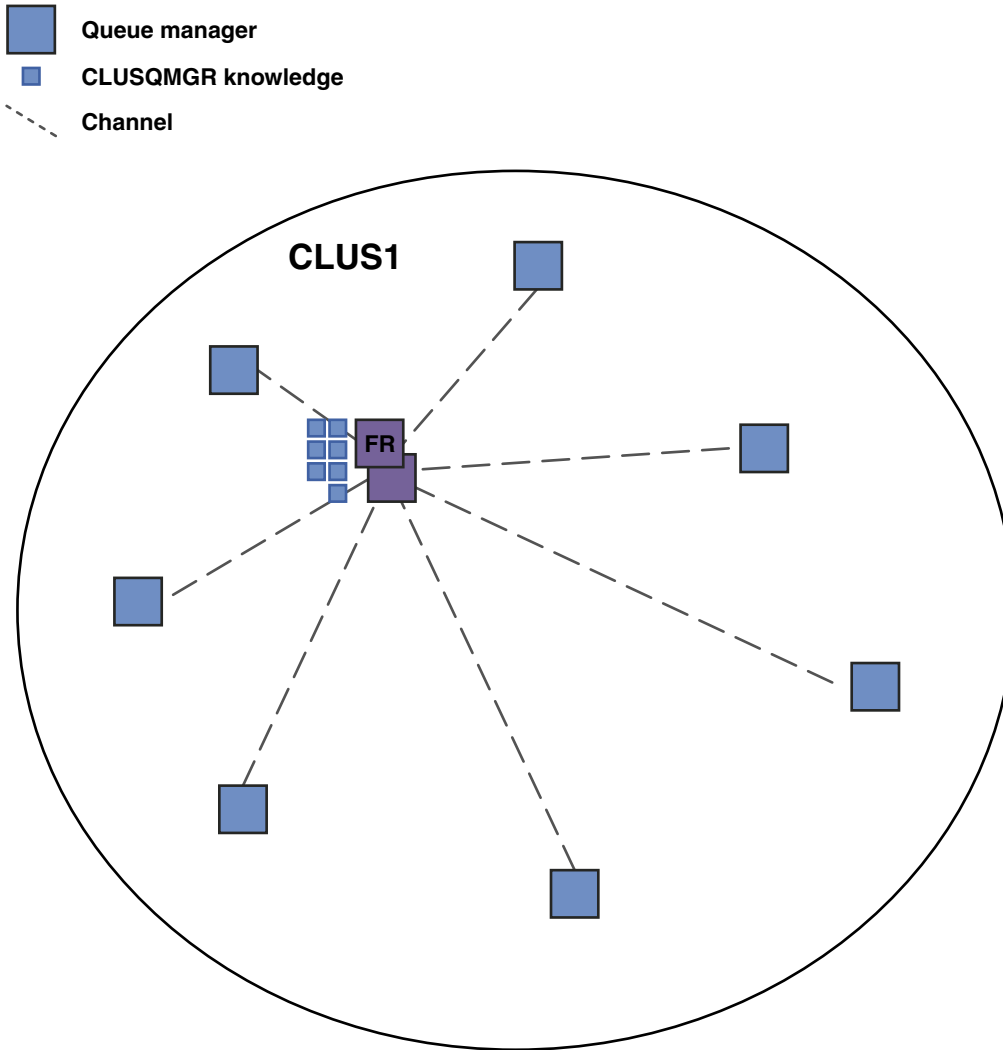


Figure 22. A queue manager cluster

In a topic host routed publish/subscribe cluster, you define the topic object on a specific queue manager in the cluster. Publish/subscribe traffic then flows through that queue manager, making it a critical queue manager in the cluster and increasing its workload. For these reasons it is not recommended to use a full repository queue manager, but to use another queue manager in the cluster. When you define the topic object on the host queue manager, knowledge of the object and its host is automatically pushed, by the full repository queue managers, to all the other queue managers in the cluster. Note that, unlike *direct routing*, each queue manager is *not* told about every other queue manager in the cluster.

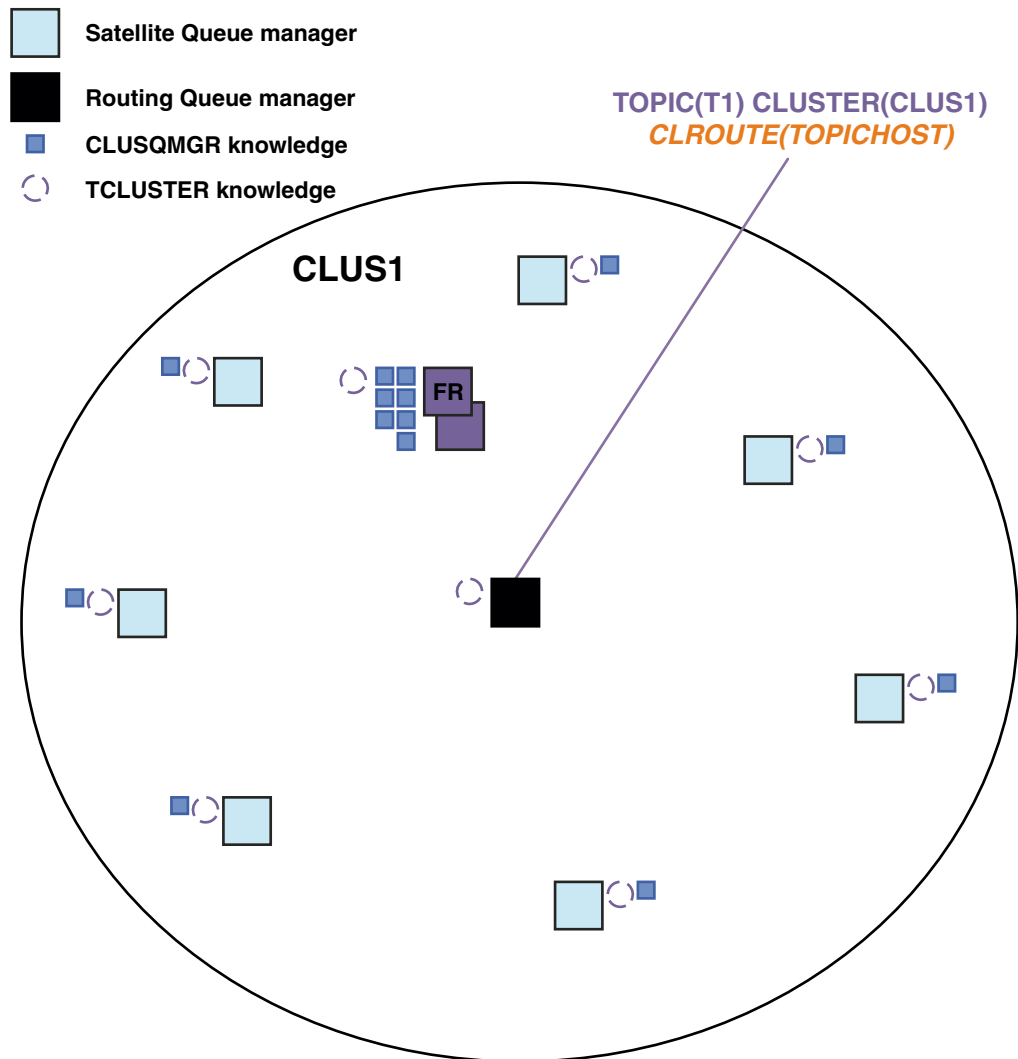


Figure 23. A topic host routed publish/subscribe cluster with one topic defined on one topic host

When a subscription is created on a queue manager, a channel is created between the subscribing queue manager and the topic host queue manager. The subscribing queue manager connects *only* to the topic host queue manager, and sends details of the subscription (in the form of a *proxy subscription*). The topic host queue manager does not forward this subscription information on to any further queue managers in the cluster.

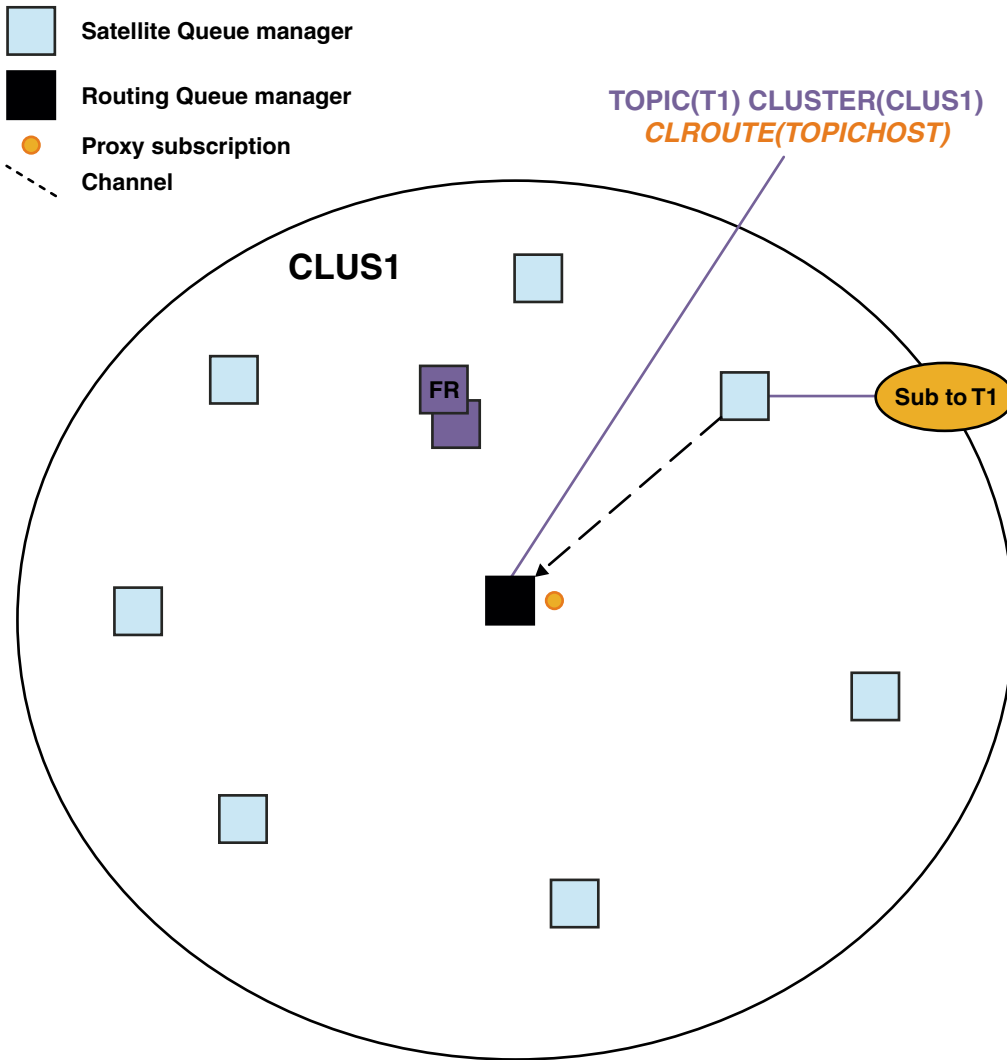


Figure 24. A topic host routed publish/subscribe cluster with one topic defined on one topic host, and one subscriber

When a publishing application connects to another queue manager and a message is published, a channel is created between the publishing queue manager and the topic host queue manager, and the message is forwarded to that queue manager. The publishing queue manager has no knowledge of any subscriptions on other queue managers in the cluster, so the message is forwarded to the topic host queue manager even if there are no subscribers to that topic in the cluster. The publishing queue manager connects *only* to the topic host queue manager. Publications are routed through the topic host to the subscribing queue managers, if any exist.

Subscriptions on the same queue manager as the publisher are satisfied directly, without first sending the messages to a topic host queue manager.

Note that, because of the critical role played by each topic host queue manager, you must choose queue managers that can handle the load, availability and connectivity requirements of topic hosting.

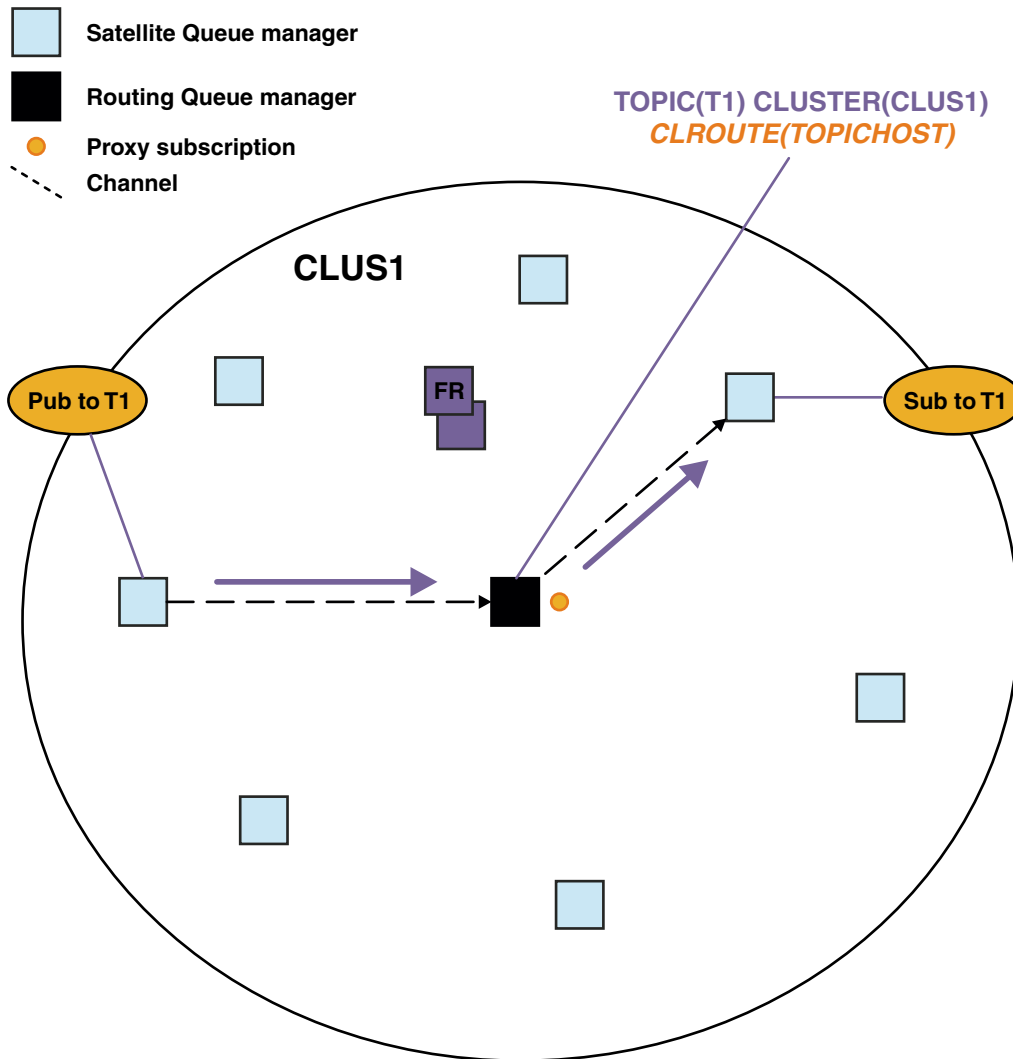


Figure 25. A topic host routed publish/subscribe cluster with one topic, one subscriber and one publisher

Dividing the topic tree across multiple queue managers

A routed topic hosting queue manager is only responsible for the subscription knowledge and publication messages that relate to the branch of the topic tree that its administered topic object is configured for. If different topics are used by different publish/subscribe applications in the cluster, you can configure different queue managers to host different clustered branches of the topic tree. This allows scaling by reducing the publication traffic, subscription knowledge and channels on each topic host queue manager in the cluster. You should use this method for distinct high volume branches of the topic tree:

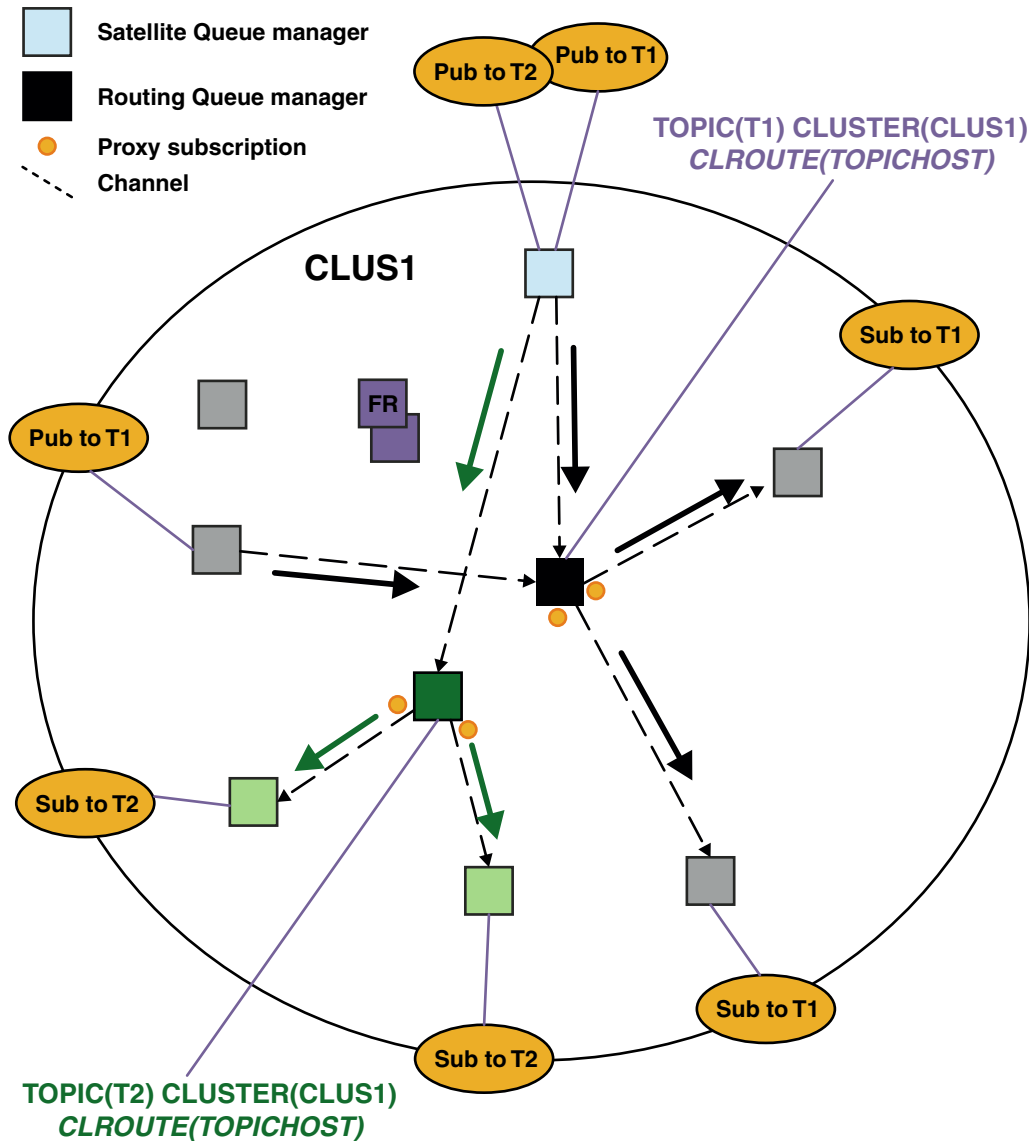


Figure 26. A topic host routed publish/subscribe cluster with two topics, each defined on one topic host

For example, using the topics described in Topic trees, if topic T1 was configured with a topic string of /USA/Alabama, and topic T2 was configured with a topic string of /USA/Alaska, then a message published to /USA/Alabama/Mobile would be routed through the queue manager hosting T1, and a message published to /USA/Alaska/Juneau would be routed through the queue manager hosting T2.

Note: You cannot make a single subscription span multiple clustered branches of the topic tree by using a wildcard higher in the topic tree than the points that are clustered. See Wildcard subscriptions.

Topic host routing using multiple topic hosts for a single topic

If a single queue manager has the responsibility for the routing of a topic, and that queue manager becomes unavailable or incapable of handling the workload, publications will not flow promptly to the subscriptions.

If you need greater resiliency, scalability and workload balancing than you get when you define a topic on just one queue manager, you can define a topic on more than one queue manager. Each individual message published is routed through a single topic host. When multiple matching topic host definitions

exist, one of the topic hosts is chosen. The choice is made in the same way as for clustered queues. This allows messages to be routed to available topic hosts, avoiding any that are unavailable, and allows message load to be workload balanced across multiple topic host queue managers and channels. However, ordering across multiple messages is not maintained when you use multiple topic hosts for the same topic in the cluster.

The following diagram shows a topic host routed cluster in which the same topic has been defined on two queue managers. In this example, the subscribing queue managers send information about the subscribed topic to both topic host queue managers in the form of a proxy subscription:

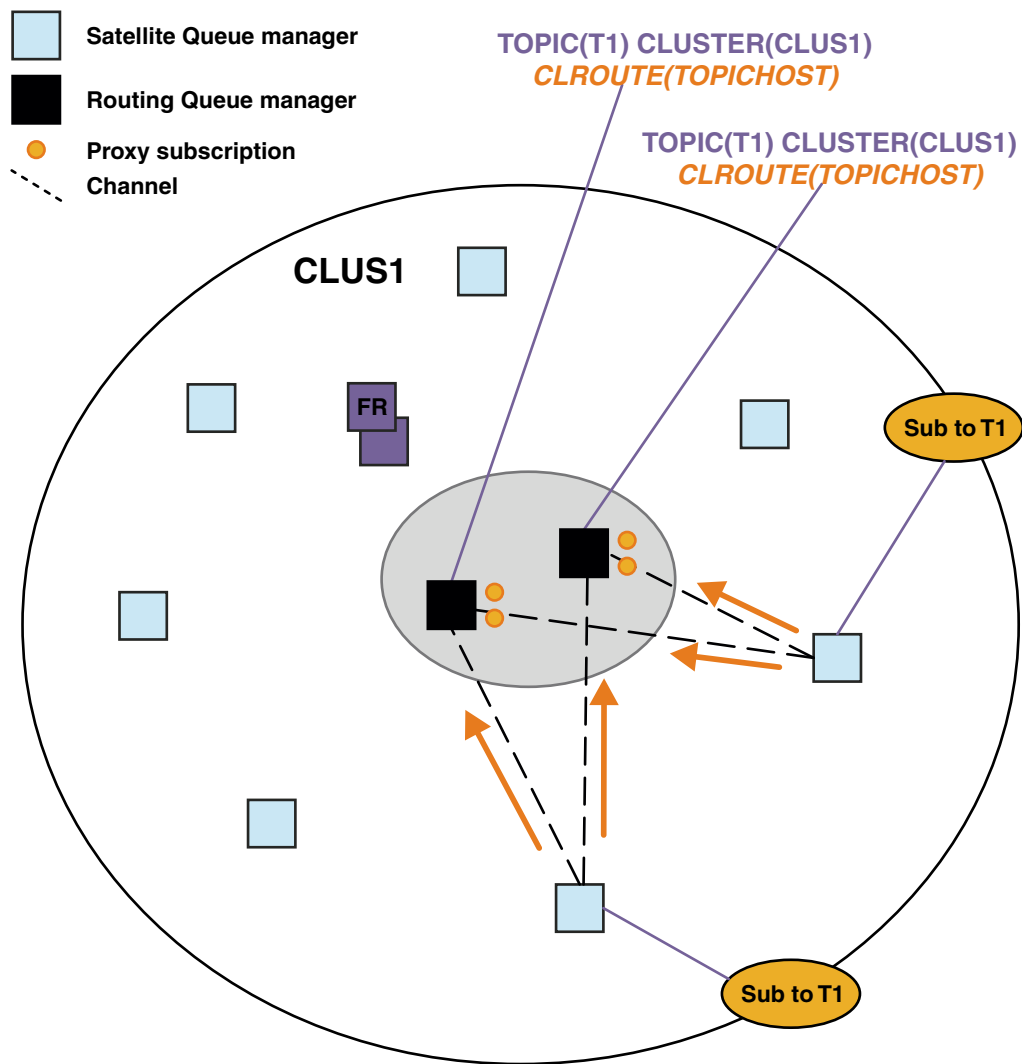


Figure 27. Creating proxy subscriptions in a multiple topic host publish/subscribe cluster

When a publication is made from a non-hosting queue manager, the queue manager sends a copy of the publication to *one* of the topic host queue managers for that topic. The system chooses the host based on the default behavior of the cluster workload management algorithm. In a typical system, this approximates to a round-robin distribution across each topic host queue manager. There is no affinity between messages from the same publishing application; this equates to using a cluster bind type of NOTFIXED.

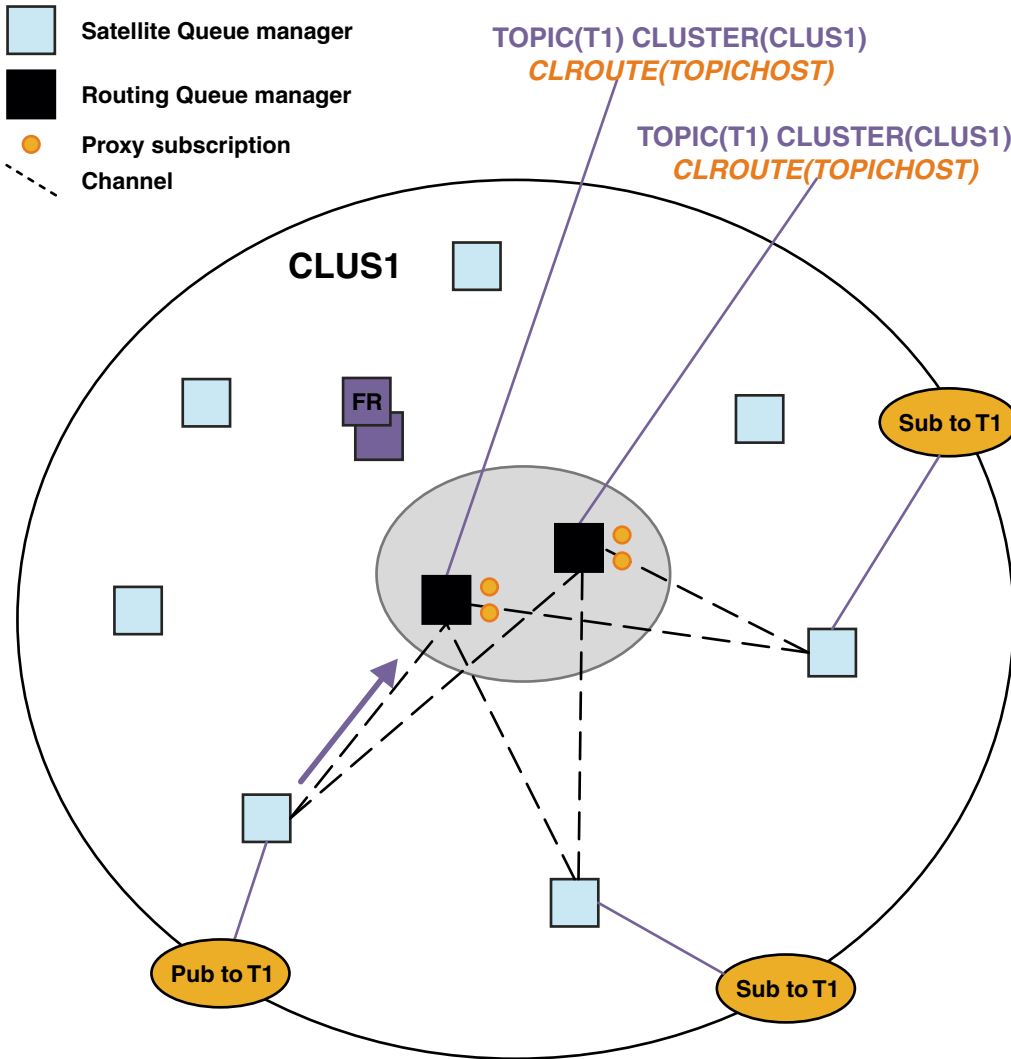


Figure 28. Receiving publications in a multiple topic host publish/subscribe cluster

Inbound publications to the chosen topic host queue manager are then forwarded to all queue managers that have registered a matching proxy subscription:

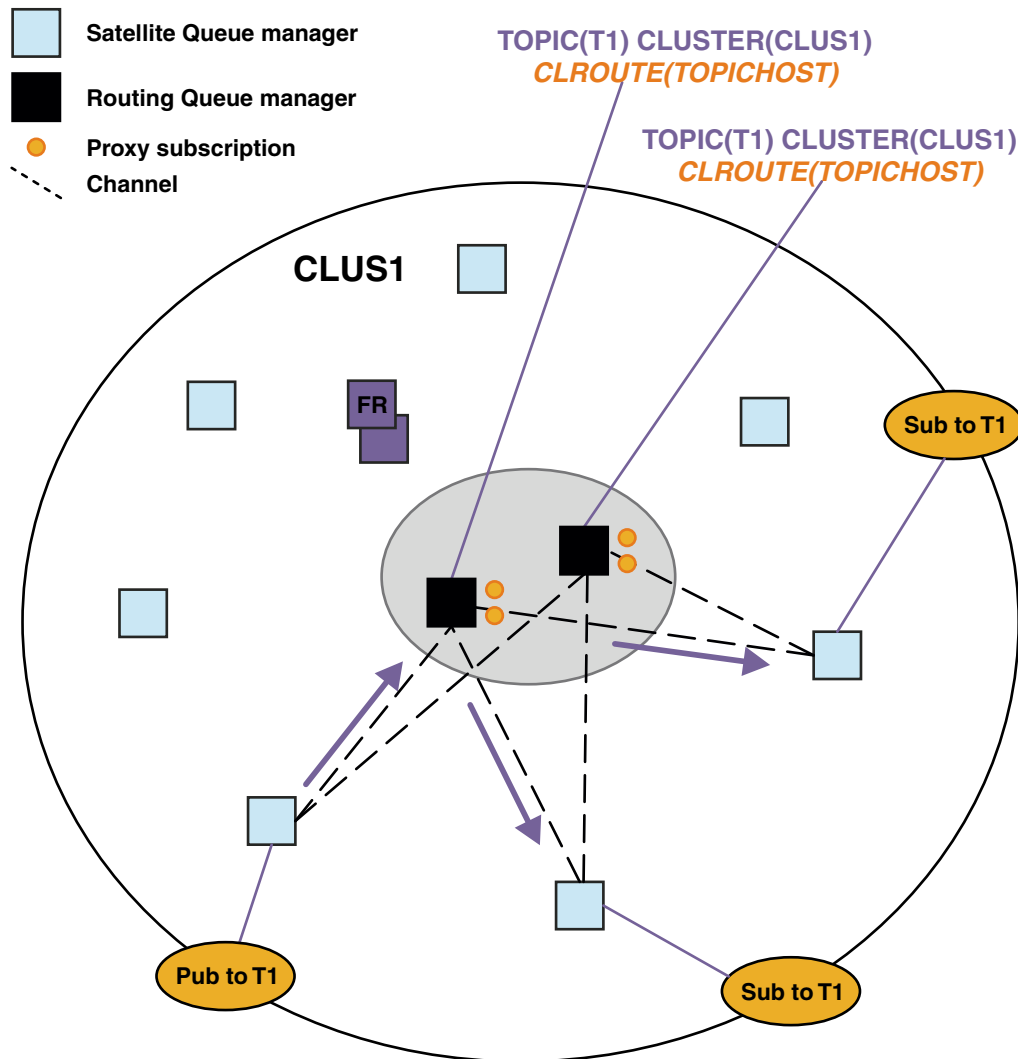


Figure 29. Routing publications to subscribers in a multiple topic host publish/subscribe cluster

Making subscriptions and publishers local to a topic host queue manager

The above examples show the routing between publishers and subscribers on queue managers that do not host administered routed topic objects. In these topologies, messages require multiple *hops* to reach the subscriptions.

Where the additional hop is not desirable, it might be appropriate to connect key publishers to topic hosting queue managers. However, if there are multiple topic hosts for a topic and only one publisher, all publication traffic will be routed through the topic host queue manager that the publisher is connected to.

Similarly, if there are key subscriptions, these could be located on a topic host queue manager. However, if there are multiple hosts of the routed topic, only a proportion of the publications will avoid the additional hop, with the remainder being routed through the other topic host queue managers first.

Topologies such as these are described further here: Topic host routing using centralized publishers or subscribers.

Note: Special planning is needed if changing the routed topic configuration when co-locating publishers or subscriptions with routed topic hosts. For example, see Adding extra topic hosts to a topic host routed cluster.

Summary and additional considerations

A topic host routed publish/subscribe cluster gives you precise control over which queue managers host each topic, and those queue managers become the *routing* queue managers for that branch of the topic tree. Moreover, queue managers without subscriptions or publishers have no need to connect with the topic host queue managers, and queue managers with subscriptions have no need to connect to queue managers that do not host a topic. This configuration can significantly reduce the number of connections between queue managers in the cluster, and the amount of information being passed between queue managers. This is especially true in large clusters where only a subset of queue managers are performing publish/subscribe work. This configuration also gives you some control over the load on individual queue managers in the cluster, so (for example) you can choose to host highly active topics on more powerful and more resilient systems. For certain configurations - notably larger clusters - it is usually a more appropriate topology than *direct routing*.

However, topic host routing also imposes certain constraints upon your system:

- System configuration and maintenance require more planning than for direct routing. You need to decide which points to cluster in the topic tree, and the location of the topic definitions in the cluster.
- Just as for direct routed topics, when a new topic host routed topic is defined, the information is pushed to the full repository queue managers, and from there direct to all members of the cluster. This event causes channels to be started to each member of the cluster from the full repositories if not already started.
- Publications are always sent to a host queue manager from a non-host queue manager, even if there are no subscriptions in the cluster. Therefore, you should use routed topics when subscriptions are typically expected to exist, or when the overhead of global connectivity and knowledge is greater than the risk of extra publication traffic.

Note: As previously described, making publishers local to a topic host can mitigate this risk.

- Messages that are published on non-host queue managers do not go direct to the queue manager that hosts the subscription, they are always routed through a topic host queue manager. This approach can increase the total overhead to the cluster, and increase message latency and reduce performance.

Note: As previously described, making subscriptions or publishers local to a topic host can mitigate this risk.

- Using a single topic host queue manager introduces a single point of failure for all messages that are published to a topic. You can remove this single point of failure by defining multiple topic hosts. However, having multiple hosts affects the order of published messages as received by subscriptions.
- Extra message load is incurred by topic host queue managers, because publication traffic from multiple queue managers needs to be processed by them. This load can be lessened: Either use multiple topic hosts for a single topic (in which case message ordering is not maintained), or use different queue managers to host routed topics for different branches of the topic tree.

Before you use topic host routing, explore the alternative approaches detailed in “Direct routing in publish/subscribe clusters” on page 63, and “Routing in publish/subscribe hierarchies” on page 93.

Publish/subscribe clustering: Best practices:

Using clustered topics makes extending the publish/subscribe domain between queue managers simple, but can lead to problems if the mechanics and implications are not fully understood. There are two models for information sharing and publication routing. Implement the model that best meets your individual business needs, and performs best on your chosen cluster.

The best practice information in the following sections does not provide a "one size fits all" solution, but rather shares common approaches to solving common problems. It assumes that you have a basic understanding of IBM MQ clusters, and of publish/subscribe messaging, and that you are familiar with the information in Distributed publish/subscribe networks and "Designing publish/subscribe clusters" on page 61.

When you use a cluster for point-to-point messaging, each queue manager in the cluster works on a "need-to-know" basis. That is, it only finds out about other cluster resources, such as other queue managers in the cluster and clustered queues, when applications connecting to them request to use them. When you add publish/subscribe messaging to a cluster, an increased level of sharing of information and connectivity between cluster queue managers is introduced. To be able to follow best practices for publish/subscribe clusters, you need to fully understand the implications of this change in behavior.

To allow you to build the best architecture, based on your precise needs, there are two models for information sharing and publication routing in publish/subscribe clusters: *direct routing* and *topic host routing*. To make the right choice, you need to understand both models, and the different requirements that each model satisfies. These requirements are discussed in the following sections, in conjunction with "Planning your distributed publish/subscribe network" on page 58:

- "Reasons to limit the number of cluster queue managers involved in publish/subscribe activity"
- "How to decide which topics to cluster" on page 79
- "How to size your system" on page 79
- "Publisher and subscription location" on page 80
- "Publication traffic" on page 81
- "Subscription change and dynamic topic strings" on page 81

Reasons to limit the number of cluster queue managers involved in publish/subscribe activity

There are capacity and performance considerations when you use publish/subscribe messaging in a cluster. Therefore, it is best practice to consider carefully the need for publish/subscribe activity across queue managers, and to limit it to only the number of queue managers that require it. After the minimum set of queue managers that need to publish and subscribe to topics are identified, they can be made members of a cluster that contains only them and no other queue managers.

This approach is especially useful if you have an established cluster already functioning well for point-to-point messaging. When you are turning an existing large cluster into a publish/subscribe cluster, it is a better practice to initially create a separate cluster for the publish/subscribe work where the applications can be tried, rather than using the current cluster. You can use a subset of existing queue managers that are already in one or more point-to-point clusters, and make this subset members of the new publish/subscribe cluster. However, the full repository queue managers for your new cluster must not be members of any other cluster; this isolates the additional load from the existing cluster full repositories.

If you cannot create a new cluster, and have to turn an existing large cluster into a publish/subscribe cluster, do not use a direct routed model. The topic host routed model usually performs better in larger clusters, because it generally restricts the publish/subscribe information sharing and connectivity to the set of queue managers that are actively performing publish/subscribe work, concentrating on the queue managers hosting the topics. The exception to that is if a manual refresh of the subscription information

is invoked on a queue manager hosting a topic definition, at which point the topic host queue manager will connect to every queue manager in the cluster. See Resynchronization of proxy subscriptions.

If you establish that a cluster cannot be used for publish/subscribe due to its size or current load, it is good practice to prevent this cluster unexpectedly being made into a publish/subscribe cluster. Use the **PSCLUS** queue manager property to stop anyone adding a clustered topic on any queue manager in the cluster. See “Inhibiting clustered publish/subscribe” on page 88.

How to decide which topics to cluster

It is important to choose carefully which topics are added to the cluster: The higher up the topic tree these topics are, the more widespread their use becomes. This can result in more subscription information and publications being propagated than necessary. If there are multiple, distinct branches of the topic tree, where some need to be clustered and some do not, create administered topic objects at the root of each branch that needs clustering and add those to the cluster. For example, if branches /A, /B and /C need clustering, define a separate clustered topic objects for each branch.

Note: The system prevents you from nesting clustered topic definitions in the topic tree. You are only permitted to cluster topics at one point in the topic tree for each sub branch. For example, you cannot define clustered topic objects for /A and for /A/B. Nesting clustered topics can lead to confusion over which clustered object applies to which subscription, especially when subscriptions are using wildcards. This is even more important when using topic host routing, where routing decisions are precisely defined by your allocation of topic hosts.

If clustered topics must be added high up the topic tree, but some branches of the tree below the clustered point do not require the clustered behavior, you can use the subscription and publication scope attributes to reduce the level of subscription and publication sharing for further topics.

You should not put the topic root node into the cluster without considering the behavior that is seen. Make global topics obvious where possible, for example by using a high-level qualifier in the topic string: /global or /cluster.

There is a further reason for not wanting to make the root topic node clustered. This is because every queue manager has a local definition for the root node, the SYSTEM.BASE.TOPIC topic object. When this object is clustered on one queue manager in the cluster, all other queue managers are made aware of it. However, when a local definition of the same object exists, its properties override the cluster object. This results in those queue managers acting as if the topic was not clustered. To resolve this, you would need to cluster every definition of SYSTEM.BASE.TOPIC. You could do this for direct routed definitions, but not for topic host routed definitions, because it causes every queue manager to become a topic host.

How to size your system

Publish/subscribe clusters typically result in a different pattern of cluster channels to point-to-point messaging in a cluster. The point-to-point model is an 'opt in' one, but publish/subscribe clusters have a more indiscriminate nature with subscription fan-out, especially when using direct routed topics. Therefore, it is important to identify which queue managers in a publish/subscribe cluster will use cluster channels to connect to other queue managers, and under what circumstances.

The following table lists the typical set of cluster sender and receiver channels expected for each queue manager in a publish/subscribe cluster under normal running, dependent on the queue manager role in the publish/subscribe cluster.

Table 5. Cluster sender and receiver channels for each routing method.

Queue manager role	Direct cluster receivers	Direct cluster senders	Topic cluster receivers	Topic cluster senders
Full repository	AllQMgrs	AllQMgrs	AllQMgrs	AllQMgrs
Host of topic definition	n/a	n/a	AllSubs+AllPubs (1)	AllSubs (1)
Subscriptions created	AllPubs (1)	AllQMgrs	AllHosts	AllHosts
Publishers connected	AllSubs (1)	AllSubs (1)	AllHosts	AllHosts
No publishers or subscribers	AllSubs (1)	None (1)	None (2)	None (2)

Key:

AllQMgrs

A channel to and from every queue manager in the cluster.

AllSubs

A channel to and from every queue manager where a subscription has been created.

AllPubs

A channel to and from every queue manager where a publishing application has been connected.

AllHosts

A channel to and from every queue manager where a definition of the clustered topic object has been configured.

None No channels to or from other queue managers in the cluster for the sole purpose of publish/subscribe messaging.

Notes:

1. If a queue manager refresh of proxy subscriptions is made from this queue manager, a channel to and from all other queue managers in the cluster might be automatically created.
2. If a queue manager refresh of proxy subscriptions is made from this queue manager, a channel to and from any other queue managers in the cluster that host a definition of a clustered topic might be automatically created.

The previous table shows that topic host routing typically uses significantly less cluster sender and receiver channels than direct routing. If channel connectivity is a concern for certain queue managers in a cluster, for reasons of capacity or ability to establish certain channels (for example, through firewalls), topic host routing is therefore a preferred solution.

Publisher and subscription location

Clustered publish/subscribe enables messages published on one queue manager to be delivered to subscriptions on any other queue manager in the cluster. As for point-to-point messaging, the cost of transmitting messages between queue managers can be detrimental to performance. Therefore you should consider creating subscriptions to topics on the same queue managers as where messages are being published.

When using topic host routing within a cluster, it is important to also consider the location of the subscriptions and publishers with respect to the topic hosting queue managers. When the publisher is not connected to a queue manager that is a host of the clustered topic, messages published are always sent to a topic hosting queue manager. Similarly, when a subscription is created on a queue manager that is not a topic host for a clustered topic, messages published from other queue managers in the cluster are always sent to a topic hosting queue manager first. More specifically, if the subscription is located on a queue

manager that hosts the topic, but there is one or more other queue managers that also host that same topic, a proportion of publications from other queue managers are routed through those other topic hosting queue managers. See Topic host routing using centralized publishers or subscribers for more information on designing a topic host routed publish/subscribe cluster to minimize the distance between publishers and subscriptions.

Publication traffic

Messages published by an application connected to one queue manager in a cluster are transmitted to subscriptions on other queue managers using cluster sender channels.

When you use direct routing, the messages published take the shortest path between queue managers. That is, they go direct from the publishing queue manager to each of the queue managers with subscriptions. Messages are not transmitted to queue managers that do not have subscriptions for the topic. See Proxy subscriptions in a publish/subscribe network.

Where the rate of publication messages between any one queue manager and another in the cluster is high, the cluster channel infrastructure between those two points must be able to maintain the rate. This might involve tuning the channels and transmission queue being used.

When you use topic host routing, each message published on a queue manager that is not a topic host is transmitted to a topic host queue manager. This is independent of whether one or more subscriptions exist anywhere else in the cluster. This introduces further factors to consider in planning:

- Is the additional latency of first sending each publication to a topic host queue manager acceptable?
- Can each topic host queue manager sustain the inbound and outbound publication rate? Consider a system with publishers on many different queue managers. If they all send their messages to a very small set of topic hosting queue managers, those topic hosts might become a bottleneck in processing those messages and routing them on to subscribing queue managers.
- Is it expected that a significant proportion of the published messages will not have a matching subscriber? If so, and the rate of publishing such messages is high, it might be best to make the publisher's queue manager a topic host. In that situation, any published message where no subscriptions exist in the cluster will not be transmitted to any other queue managers.

These problems might also be eased by introducing multiple topic hosts, to spread the publication load across them:

- Where there are multiple distinct topics, each with a proportion of the publication traffic, consider hosting them on different queue managers.
- If the topics cannot be separated onto different topic hosts, consider defining the same topic object on multiple queue managers. This results in publications being workload balanced across each of them for routing. However, this is only appropriate when publication message ordering is not required.

Subscription change and dynamic topic strings

Another consideration is the effect on performance of the system for propagating proxy subscriptions. Typically, a queue manager sends a proxy subscription message to certain other queue managers in the cluster when the first subscription for a specific clustered topic string (not just a configured topic object) is created on that queue manager. Similarly, a proxy subscription deletion message is sent when the last subscription for a specific clustered topic string is deleted.

For direct routing, each queue manager with subscriptions sends those proxy subscriptions to every other queue manager in the cluster. For topic host routing, each queue manager with subscriptions only sends the proxy subscriptions to each queue manager that hosts a definition for that clustered topic. Therefore, with direct routing, the more queue managers there are in the cluster, the higher the overhead of maintaining proxy subscriptions across them. Whereas, with topic host routing, the number of queue managers in the cluster is not a factor.

In both routing models, if a publish/subscribe solution consists of many unique topic strings being subscribed to, or the topics on a queue manager in the cluster are frequently being subscribed and unsubscribed, a significant overhead will be seen on that queue manager, caused by constantly generating messages distributing and deleting the proxy subscriptions. With direct routing, this is compounded by the need to send these messages to every queue manager in the cluster.

If the rate of change of subscriptions is too high to accommodate, even within a topic host routed system, see *Subscription performance in publish/subscribe networks* for information about ways to reduce proxy subscription overhead.

Defining cluster topics:

Cluster topics are administrative topics with the **cluster** attribute defined. Information about cluster topics is pushed to all members of a cluster, and combined with local topics to create portions of a topic space that spans multiple queue managers. This enables messages published on a topic on one queue manager to be delivered to subscriptions of other queue managers in the cluster.

When you define a cluster topic on a queue manager, the cluster topic definition is sent to the full repository queue managers. The full repositories then propagate the cluster topic definition to all queue managers within the cluster, making the same cluster topic available to publishers and subscribers at any queue manager in the cluster. The queue manager on which you create a cluster topic is known as a cluster topic host. The cluster topic can be used by any queue manager in the cluster, but any modifications to a cluster topic must be made on the queue manager where that topic is defined (the host) at which point the modification is propagated to all members of the cluster through the full repositories.

When you use direct routing, the location of the clustered topic definition does not directly affect the behavior of the system, because all queue managers in the cluster use the topic definition in the same way. You should therefore define the topic on any queue manager that will be a member of the cluster for as long as the topic is needed, and that is on a system reliable enough to regularly be in contact with the full repository queue managers.

When you use topic host routing, the location of the clustered topic definition is very important, because other queue managers in the cluster create channels to this queue manager and send subscription information and publications to it. To choose the best queue manager to host the topic definition, you need to understand topic host routing. See “Topic host routing in publish/subscribe clusters” on page 68.

If you have a clustered topic, and a local topic object, then the local topic takes precedence. See “Multiple cluster topic definitions of the same name” on page 85.

For information about the commands to use to display cluster topics, see the related information.

Clustered topic inheritance

Typically, publishing and subscribing applications in a clustered publish/subscribe topology expect to work the same, no matter which queue manager in the cluster they are connected to. This is why clustered administered topic objects are propagated to every queue manager in the cluster.

An administered topic object inherits its behavior from other administered topic objects higher in the topic tree. This inheritance occurs when an explicit value has not been set for a topic parameter.

In the case of clustered publish/subscribe, it is important to consider such inheritance because it introduces the possibility that publishers and subscribers will behave differently depending on which queue manager they connect to. If a clustered topic object leaves any parameters to inherit from higher topic objects, the topic might behave differently on different queue managers in the cluster. Similarly, locally defined topic objects defined below a clustered topic object in the topic tree will mean those lower

topics are still clustered, but the local object might change its behavior in some way that differs from other queue managers in the cluster.

Wildcard subscriptions

Proxy subscriptions are created when local subscriptions are made to a topic string that resolves at, or below, a clustered topic object. If a wildcard subscription is made higher in the topic hierarchy than any cluster topic, it does not have proxy subscriptions sent around the cluster for the matching cluster topic, and therefore receives no publications from other members of the cluster. It does however receive publications from the local queue manager.

However, if another application subscribes to a topic string that resolves to or below the cluster topic, proxy subscriptions are generated and publications are propagated to this queue manager. On arrival the original, higher wildcard subscription is considered a legitimate recipient of those publications and receives a copy. If this behavior is not required, set **WILDCARD(BLOCK)** on the clustered topic. This makes the original wildcard not be considered a legitimate subscription, and stops it receiving any publications (local, or from elsewhere in the cluster) on the cluster topic, or its subtopics.

Related information:

Working with administrative topics

Working with subscriptions

DISPLAY TOPIC

DISPLAY TPSTATUS

DISPLAY SUB

Cluster topic attributes:

When a topic object has the cluster name attribute set, the topic definition is propagated across all queue managers in the cluster. Each queue manager uses the propagated topic attributes to control the behavior of publish/subscribe applications.

A topic object has a number of attributes that apply to publish/subscribe clusters. Some control the general behavior of the publishing and subscribing applications and some control how the topic is used across the cluster.

A clustered topic object definition must be configured in a way that all queue managers in the cluster can correctly use it.

For example if the model queues to be used for managed subscriptions (MDURMDL and MNDURMDL) are set to a non-default queue name, that named model queue must be defined on all queue managers where managed subscriptions will be created.

Similarly, if any attribute is set to ASPARENT, the behavior of the topic will be dependent on the higher nodes in the topic tree (see Administrative topic objects) on each individual queue manager in the cluster. This might result in different behavior when publishing or subscribing from different queue managers.

The main attributes that directly relate to publish/subscribe behavior across the cluster are as follows:

CLROUTE

This parameter controls the routing of messages between queue managers where publishers are connected, and queue managers where matching subscriptions exist.

- You configure the route to be either direct between these queue managers, or through a queue manager that hosts a definition of the clustered topic. See Publish/subscribe clusters for more details.

- You cannot change the **CLROUTE** while the **CLUSTER** parameter is set. To change the **CLROUTE**, first set the **CLUSTER** property to be blank. This stops applications that use the topic from behaving in a clustered manner. This in turn results in a break in publications being delivered to subscriptions, so you should also quiesce publish/subscribe messaging while making the change.

PROXYSUB

This parameter controls when proxy subscriptions are made.

- **FIRSTUSE** is the default value, and causes proxy subscriptions to be sent in response to local subscriptions on a queue manager in a distributed publish/subscribe topology, and canceled when no longer required. For details about why you might want to change this attribute from the default value of **FIRSTUSE**, see *Individual proxy subscription forwarding and **publish everywhere***.
- To enable *publish everywhere*, you set the **PROXYSUB** parameter to **FORCE** for a high-level topic object. This results in a single wildcard proxy subscription that matches all topics below this topic object in the topic tree.

Note: Setting the **PROXYSUB(FORCE)** attribute in a large or busy publish/subscribe cluster can result in excessive load on system resources. The **PROXYSUB(FORCE)** attribute is propagated to every queue manager, not just the queue manager that the topic was defined on. This causes every queue manager in the cluster to create a wildcarded proxy subscription.

A copy of a message to this topic, published on any queue manager in the cluster, is sent to every queue manager in the cluster - either directly, or through a topic host queue manager, depending on the **CLROUTE** setting.

When the topic is direct routed, every queue manager creates cluster sender channels to every other queue manager. When the topic is topic host routed, channels to each topic host queue manager are created from every queue manager in the cluster.

For more information about the **PROXYSUB** parameter when used in clusters, see *Direct routed publish/subscribe performance*.

PUBSCOPE and SUBSCOPE

These parameters determine whether this queue manager propagates publications to queue managers in the topology (publish/subscribe cluster or hierarchy) or restricts the scope to just its local queue manager. You can do the equivalent job programmatically using **MQPMO_SCOPE_QMGR** and **MQSO_SCOPE_QMGR**.

PUBSCOPE

If a cluster topic object is defined with **PUBSCOPE(QMGR)**, the definition is shared with the cluster, but the scope of publications that are based on that topic is local only and they are not sent to other queue managers in the cluster.

SUBSCOPE

If a cluster topic object is defined with **SUBSCOPE(QMGR)**, the definition is shared with the cluster, but the scope of subscriptions that are based on that topic is local only, therefore no proxy subscriptions are sent to other queue managers in the cluster.

These two attributes are commonly used together to isolate a queue manager from interacting with other members of the cluster on particular topics. The queue manager neither publishes or receives publications on those topics to and from other members of the cluster. This situation does not prevent publication or subscription if topic objects are defined on subtopics.

Setting **SUBSCOPE** to **QMGR** on a local definition of a topic does not prevent other queue managers in the cluster from propagating their proxy subscriptions to the queue manager if they are using a clustered version of the topic, with **SUBSCOPE(ALL)**. However, if the local definition also sets **PUBSCOPE** to **QMGR** those proxy subscriptions are not sent publications from this queue manager.

Related information:

Publication scope

Subscription scope




Multiple cluster topic definitions of the same name:

You can define the same named cluster topic object on more than one queue manager in the cluster, and in certain scenarios this enables specific behavior. When multiple cluster topic definitions exist of the same name, the majority of properties should match. If they do not, errors or warnings are reported depending on the significance of the mismatch.

In general, if there is a mismatch in the properties of multiple cluster topic definitions, warnings are issued and one of the topic object definitions is used by each queue manager in the cluster. Which definition is used by each queue manager is not deterministic, or consistent across the queue managers in the cluster. Such mismatches should be resolved as quickly as possible.

During cluster setup or maintenance you sometimes need to create multiple cluster topic definitions that are not identical. However this is only ever useful as a temporary measure, and it is therefore treated as a potential error condition.




When mismatches are detected, the following warning messages are written to each queue manager's error log:

-   On distributed systems and IBM i, AMQ9465 and AMQ9466.
-  On z/OS systems, CSQX465I and CSQX466I.

The chosen properties for any topic string on each queue manager can be determined by viewing topic status rather than the topic object definitions, for example by using **DISPLAY TPSTATUS**.

In some situations, a conflict in configuration properties is severe enough to stop the topic object being created, or to cause the mismatched objects to be marked as invalid and not propagated across the cluster (See **CLSTATE** in **DISPLAY TOPIC**). These situations occur when there is a conflict in the cluster routing property (**CLROUTE**) of the topic definitions. Additionally, due to the importance of consistency across topic host routed definitions, further inconsistencies are rejected as detailed in subsequent sections of this article.

If the conflict is detected at the time that the object is defined, the configuration change is rejected. If detected later by the full repository queue managers, the following warning messages are written to the queue managers error logs:

-   On distributed systems and IBM i, AMQ9879
-  On z/OS systems, CSQX879E.

When multiple definitions of the same topic object are defined in the cluster, a locally defined definition takes precedence over any remotely defined one. Therefore, if any differences exist in the definitions, the queue managers hosting the multiple definitions behave differently from each other.

The effect of defining a non-cluster topic with the same name as a cluster topic from another queue manager

It is possible to define an administered topic object that is not clustered on a queue manager that is in a cluster, and simultaneously define the same named topic object as a clustered topic definition on a different queue manager. In this case, the locally defined topic object takes precedence over all remote definitions of the same name.

This has the effect of preventing the clustering behavior of the topic when used from this queue manager. That is, subscriptions might not receive publications from remote publishers, and messages from publishers might not be propagated to remote subscriptions in the cluster.

Careful consideration should be given before configuring such a system, because this can lead to confusing behavior.

Note: If an individual queue manager needs to prevent publications and subscriptions from propagating around the cluster, even when the topic has been clustered elsewhere, an alternative approach is to set the publication and subscription scopes to only the local queue manager. See “Cluster topic attributes” on page 83.

Multiple cluster topic definitions in a direct routed cluster

For direct routing, you do not usually define the same cluster topic on more than one cluster queue manager. This is because direct routing makes the topic available at all queue managers in the cluster, no matter which queue manager it was defined on. Moreover, adding multiple cluster topic definitions significantly increases system activity and administrative complexity, and with increased complexity comes a greater chance of human error:

- Each definition results in an additional cluster topic object being pushed out to the other queue managers in the cluster, including the other cluster topic host queue managers.
- All definitions for a specific topic in a cluster must be identical, otherwise it is difficult to work out which topic definition is being used by a queue manager.

It is also not essential that the sole host queue manager is continually available for the topic to function correctly across the cluster, because the cluster topic definition is cached by the full repository queue managers and by all other queue managers in their partial cluster repositories. For more information, see Availability of topic host queue managers that use direct routing.

For a situation in which you might need to temporarily define a cluster topic on a second queue manager, for example when the existing host of the topic is to be removed from the cluster, see Moving a cluster topic definition to a different queue manager.

If you need to alter a cluster topic definition, take care to modify it at the same queue manager it was defined on. Attempting to modify it from another queue manager might accidentally create a second definition of the topic with conflicting topic attributes.

Multiple cluster topic definitions in a topic host routed cluster

When a cluster topic is defined with a cluster route of *topic host*, the topic is propagated across all queue managers in the cluster just as for *direct* routed topics. Additionally, all publish/subscribe messaging for that topic is routed through the queue managers where that topic is defined. Therefore the location and number of definitions of the topic in the cluster becomes important (see “Topic host routing in publish/subscribe clusters” on page 68).

To ensure adequate availability and scalability it is useful, if possible, to have multiple topic definitions. See Availability of topic host queue managers that use topic host routing.

When adding or removing additional definitions of a *topic host* routed topic in a cluster, you should consider the flow of messages at the time of the configuration change. If messages are being published in the cluster to the topic at the time of the change, a staged process is required to add or remove a topic definition. See Moving a cluster topic definition to a different queue manager and Adding extra topic hosts to a topic host routed cluster.

As previously explained, the properties of the multiple definitions should match, with the possible exception of the **PUB** parameter, as described in the next section. When publications are routed through topic host queue managers it is even more important for multiple definitions to be consistent. Therefore, an inconsistency detected in either the topic string or cluster name is rejected if one or more of the topic definitions has been configured for topic host cluster routing.

Note: Cluster topic definitions are also rejected if an attempt is made to configure them above or below another topic in the topic tree, where the existing clustered topic definition is configured for topic host routing. This prevents ambiguity in the routing of publications with respect to wildcarded subscriptions.

Special handling for the PUB parameter

The **PUB** parameter is used to control when applications can publish to a topic. In the case of topic host routing in a cluster, it can also control which topic host queue managers are used to route publications. For this reason it is permitted to have multiple definitions of the same topic object in the cluster, with different settings for the **PUB** parameter.

If multiple remote clustered definitions of a topic have different settings for this parameter, the topic allows publications to be sent and delivered to subscriptions if the following conditions are met:

- There is not a matching topic object defined on the queue manager that the publisher is connected to that is set to **PUB(DISABLED)**.
- One or more of the multiple topic definitions in the cluster is set to **PUB(ENABLED)**, or one or more of the multiple topic definitions is set to **PUB(ASPARENT)** and the local queue managers where the publisher is connected and the subscription defined are set to **PUB(ENABLED)** at a higher point in the topic tree.

For topic host routing, when messages are published by applications connected to queue managers that are not topic hosts, messages are only routed to the topic hosting queue managers where the **PUB** parameter has not explicitly been set to **DISABLED**. You can therefore use the **PUB(DISABLED)** setting to quiesce message traffic through certain topic hosts. You might want to do this to prepare for maintenance or removal of a queue manager, or for the reasons described in Adding extra topic hosts to a topic host routed cluster.

Availability of cluster topic host queue managers:

Design your publish/subscribe cluster to minimize the risk that, should a topic host queue manager become unavailable, the cluster will no longer be able to process traffic for the topic. The effect of a topic host queue manager becoming unavailable depends on whether the cluster is using topic host routing or direct routing.

Availability of topic host queue managers that use direct routing

For direct routing, you do not usually define the same cluster topic on more than one cluster queue manager. This is because direct routing makes the topic available at all queue managers in the cluster, no matter which queue manager it was defined on. See Multiple cluster topic definitions in a direct routed cluster.

In a cluster, whenever the host of a clustered object (for example a clustered queue or clustered topic) becomes unavailable for a prolonged period of time, the other members of the cluster will eventually expire the knowledge of those objects. In the case of a clustered topic, if the cluster topic host queue manager becomes unavailable, the other queue managers continue to process publish/subscribe requests for the topic in a direct clustered way (that is, sending publications to subscriptions on remote queue managers) for at least 60 days from when the topic hosting queue manager was last in communication with the full repository queue managers. If the queue manager on which you defined the cluster topic object is never made available again, then eventually the cached topic objects on the other queue

managers are deleted and the topic reverts to a local topic, in which case subscriptions cease to receive publications from applications connected to remote queue managers.

With the 60 day period to recover the queue manager on which you define a cluster topic object, there is little need to take special measures to guarantee that a cluster topic host remains available (note, however, that any subscriptions defined on the unavailable cluster topic host do not remain available). The 60 day period is sufficient to cater for technical problems, and is likely to be exceeded only because of administrative errors. To mitigate that possibility, if the cluster topic host is unavailable, all members of the cluster write error log messages hourly, stating that their cached cluster topic object was not refreshed. Respond to these messages by making sure that the queue manager on which the cluster topic object is defined, is running. If it is not possible to make the cluster topic host queue manager available again, define the same clustered topic definition, with exactly the same attributes, on another queue manager in the cluster.

Availability of topic host queue managers that use topic host routing

For topic host routing, all publish/subscribe messaging for a topic is routed through the queue managers where that topic is defined. For this reason, it is very important to consider the continual availability of these queue managers in the cluster. If a topic host becomes unavailable, and no other host exists for the topic, traffic from publishers to subscribers on different queue managers in the cluster immediately halts for the topic. If additional topic hosts are available, the cluster queue managers route new publication traffic through these topic hosts, providing continuous availability of message routes.

As for direct topics, after 60 days, if the first topic host is still unavailable, knowledge of that topic host's topic is removed from the cluster. If this is the last remaining definition for this topic in the cluster, all other queue managers cease to forward publications to any topic host for routing.

To ensure adequate availability and scalability it is therefore useful, if possible, to define each topic on at least two cluster queue managers. This gives protection against any given topic host queue manager becoming unavailable. See also Multiple cluster topic definitions in a topic host routed cluster.

If you cannot configure multiple topic hosts (for example because you need to preserve message ordering), and you cannot configure just one topic host (because the availability of a single queue manager must not affect the flow of publications to subscriptions across all queue managers in the cluster), consider configuring the topic as a direct routed topic. This avoids reliance on a single queue manager for the whole cluster, but does still require each individual queue manager to be available in order for it to process locally hosted subscriptions and publishers.

Inhibiting clustered publish/subscribe:

Introducing the first direct routed clustered topic into a cluster forces every queue manager in the cluster to become aware of every other queue manager, and potentially causes them to create channels to each other. If this is not desirable, you should instead configure topic host routed publish/subscribe. If the existence of a direct routed clustered topic might jeopardize the stability of the cluster, due to scaling concerns of each queue manager, you can completely disable clustered publish/subscribe functionality by setting **PSCLUS** to **DISABLED** on every queue manager in the cluster.

As described in "Direct routing in publish/subscribe clusters" on page 63, when you introduce a direct routed clustered topic into a cluster, all partial repositories are automatically notified of all other members of the cluster. The clustered topic might also create subscriptions at all other nodes (for example, where **PROXYSUB(FORCE)** is specified) and cause large numbers of channels to be started from a queue manager, even when there are no local subscriptions. This puts an immediate additional load on each queue manager in the cluster. For a cluster that contains many queue managers, this can cause a significant reduction in performance. Therefore the introduction of direct routed publish/subscribe to a cluster must be carefully planned.

When you know that a cluster cannot accommodate the overheads of direct routed publish/subscribe, you can instead use topic host routed publish/subscribe. For an overview of the differences, see “Designing publish/subscribe clusters” on page 61.

If you prefer to completely disable publish/subscribe functionality for the cluster, you can do so by setting the queue manager attribute **PSCLUS** to DISABLED on every queue manager in the cluster. This setting disables both direct routed and topic host routed publish/subscribe in the cluster, by modifying three aspects of queue manager functionality:

- An administrator of this queue manager is no longer able to define a Topic object as clustered.
- Incoming topic definitions or proxy subscriptions from other queue managers are rejected, and a warning message is logged to inform the administrator of incorrect configuration.
- Full repositories no longer automatically share information about every queue manager with all other partial repositories when they receive a topic definition.

Although **PSCLUS** is a parameter of each individual queue manager in a cluster, it is not intended to selectively disable publish/subscribe in a subset of queue managers in the cluster. If you selectively disable in this way, you will see frequent error messages. This is because proxy subscriptions and topic definitions are constantly seen and rejected if a topic is clustered on a queue manager where **PSCLUS** is enabled.

You should therefore aim to set **PSCLUS** to DISABLED on every queue manager in the cluster. However, in practice this state can be difficult to achieve and maintain, for example queue managers can join and leave the cluster at any time. At the very least, you must ensure that **PSCLUS** is set to DISABLED on all full repository queue managers. If you do this, and a clustered topic is subsequently defined on an ENABLED queue manager in the cluster, this does not cause the full repositories to inform every queue manager of every other queue manager, and so your cluster is protected from potential scaling issues across all queue managers. In this scenario, the origin of the clustered topic is reported in the error logs of the full repository queue managers.

If a queue manager participates in one or more publish/subscribe clusters, and also one or more point-to-point clusters, you must set **PSCLUS** to ENABLED on that queue manager. For this reason, when overlapping a point-to-point cluster with a publish subscribe cluster, you should use a separate set of full repositories in each cluster. This approach allows topic definitions and 'all queue manager' information to flow only in the publish/subscribe cluster.

To avoid inconsistent configurations when you change **PSCLUS** from ENABLED to DISABLED, no clustered topic objects can exist in any cluster of which this queue manager is a member. Any such topics, even remotely defined ones, must be deleted before changing **PSCLUS** to DISABLED.

For more information about **PSCLUS**, see ALTER QMGR (PSCLUS).

Related information:

Direct routed publish/subscribe cluster performance

Publish/subscribe and multiple clusters:

A single queue manager can be a member of more than one cluster. This arrangement is sometimes known as *overlapping clusters*. Through such an overlap, clustered queues can be made accessible from multiple clusters, and point-to-point message traffic can be routed from queue managers in one cluster to queue managers in another cluster. Clustered topics in publish/subscribe clusters do not provide the same capability. Therefore, their behavior must be clearly understood when using multiple clusters.

Unlike for a queue, you cannot associate a topic definition with more than one cluster. The scope of a clustered topic is confined to those queue managers in the same cluster as the topic is defined for. This allows publications to be propagated to subscriptions only on those queue managers in the same cluster.

A queue manager's topic tree




Figure 30. Overlapping clusters: Two clusters each subscribing to different topics

When a queue manager is a member of multiple clusters it is made aware of all clustered topics defined in each of those clusters. For example, in the previous figure QM3 is aware of both the T_B and T_C administered clustered topic objects, whereas QM1 is only aware of T_B . QM3 applies both topic definitions to its local topic, and therefore has a different behavior to QM1 for certain topics. For this reason it is important that clustered topics from different clusters do not interfere with each other. Interference can occur when one clustered topic is defined above or below another clustered topic in a different cluster (for example, they have topic strings of /Sport and /Sport/Football) or even for the same topic string in both. Another form of interference is when administered clustered topic objects are defined with the same object name in different clusters, but for different topic strings.

If such a configuration is made, the delivery of publications to matching subscriptions becomes very dependent on the relative locations of publishers and subscribers with respect to the cluster. For this reason, you cannot rely on such a configuration, and you should change it to remove the interfering topics.

When planning an overlapping cluster topology with publish/subscribe messaging, you can avoid any interference by treating the topic tree and clustered topic object names as if they span all overlapping clusters in the topology.

Integrating multiple publish/subscribe clusters

If there is a requirement for publish/subscribe messaging to span queue managers in different clusters, there are two options available:

- Connect the clusters together through the use of a publish/subscribe hierarchy configuration. See *Combining the topic spaces of multiple clusters*.
- Create an additional cluster that overlays the existing clusters and includes all queue managers that need to publish or subscribe to a particular topic.

With the latter option, you should consider carefully the size of the cluster and the most effective cluster routing mechanism. See “*Designing publish/subscribe clusters*” on page 61.

Design considerations for retained publications in publish/subscribe clusters:

There are a few restrictions to consider when designing a publish/subscribe cluster to work with retained publications.

Considerations

Consideration 1: The following cluster queue managers always store the latest version of a retained publication:

- The publisher's queue manager
- In a topic host routed cluster, the topic host (provided there is only one topic host for the topic, as explained in the next section of this article)
- All queue managers with subscriptions matching the topic string of the retained publication

Consideration 2: Queue managers do not receive updated retained publications while they have no subscriptions. Therefore any retained publication stored on a queue manager that no longer subscribes to the topic will become stale.

Consideration 3: On creating any subscription, if there is a local copy of a retained publication for the topic string, the local copy is delivered to the subscription. If you are the first subscriber for any given topic string, a matching retained publication is also delivered from one of the following cluster members:

- In a direct routed cluster, the publisher's queue manager
- In a topic host routed cluster, the topic hosts for the given topic

Delivery of a retained publication from a topic host or publishing queue manager to the subscribing queue manager is asynchronous to the MQSUB calls. Therefore, if you use the MQSUBRQ call, the latest retained publication might be missed until a subsequent call to MQSUBRQ.

Implications

For any publish/subscribe cluster, when a first subscription is made, the local queue manager might be storing a stale copy of a retained publication and this is the copy that is delivered to the new subscription. The existence of a subscription on the local queue manager means that this will resolve the next time the retained publication is updated.

For a topic host routed publish/subscribe cluster, if you configure more than one topic host for a given topic, new subscribers might receive the latest retained publication from a topic host, or they might receive a stale retained publication from another topic host (with the latest having been lost). For topic host routing, it is usual to configure multiple topic hosts for a given topic. However, if you expect applications to make use of retained publications, you should configure only one topic host for each topic.

For any given topic string, you should use only a single publisher, and ensure the publisher always uses the same queue manager. If you do not do this, different retained publications might be active at different queue managers for the same topic, leading to unexpected behavior. Because multiple proxy subscriptions are distributed, multiple retained publications might be received.

If you are still concerned about subscribers using stale publications, consider setting a message expiry when you create each retained publication.

You can use the **CLEAR TOPICSTR** command to remove a retained publication from a publish/subscribe cluster. In certain circumstances you might need to issue the command on multiple members of the publish/subscribe cluster, as described in **CLEAR TOPICSTR** .

Wildcard subscriptions and retained publications

If you are using wildcard subscriptions, the corresponding proxy subscriptions delivered to other members of the publish/subscribe cluster are wildcarded from the topic separator immediately prior to the first wildcard character. See Wildcards and cluster topics.

Therefore the wildcard used might match more topic strings, and more retained publications, than will match the subscribing application.

This increases the amount of storage needed for the retained publications, and you therefore need to ensure that the hosting queue managers have enough storage capacity.

Related information:

Retained publications

Individual proxy subscription forwarding and publish everywhere

REFRESH CLUSTER considerations for publish/subscribe clusters:

Issuing the **REFRESH CLUSTER** command results in the queue manager temporarily discarding locally held information about a cluster, including any cluster topics and their associated proxy subscriptions.

The time taken from issuing the **REFRESH CLUSTER** command to the point that the queue manager regains a full knowledge of the necessary information for clustered publish/subscribe depends on the size of the cluster, the availability, and the responsiveness of the full repository queue managers.

During the refresh processing, disruption to publish/subscribe traffic in a publish/subscribe cluster occurs. For large clusters, use of the **REFRESH CLUSTER** command can disrupt the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster. For these reasons, the **REFRESH CLUSTER** command must be used in a publish/subscribe cluster only when under the guidance of your IBM Support Center.

The disruption to the cluster can appear externally as the following symptoms:

- Subscriptions to cluster topics on this queue manager are not receiving publications from publishers that are connected to other queue managers in the cluster.
- Messages that are published to cluster topics on this queue manager are not being propagated to subscriptions on other queue managers.
- Subscriptions to cluster topics on this queue manager created during this period are not consistently sending proxy subscriptions to other members of the cluster.
- Subscriptions to cluster topics on this queue manager deleted during this period are not consistently removing proxy subscriptions from other members of the cluster.
- 10-second pauses, or longer, in message delivery.
- **MQPUT** failures, for example, **MQRC_PUBLICATION_FAILURE**.

- Publications placed on the dead-letter queue with a reason of MQRC_UNKNOWN_REMOTE_Q_MGR
- For these reasons publish/subscribe applications need to be quiesced before issuing the **REFRESH CLUSTER** command.

See also Usage notes for **REFRESH CLUSTER** and “Clustering: Using REFRESH CLUSTER best practices” on page 55.

After a **REFRESH CLUSTER** command is issued on a queue manager in a publish/subscribe cluster, wait until all cluster queue managers and cluster topics have been successfully refreshed, then resynchronize proxy subscriptions as described in Resynchronization of proxy subscriptions. When all proxy subscriptions have been correctly resynchronized, restart your publish/subscribe applications.

If a **REFRESH CLUSTER** command is taking a long time to complete, monitor it by looking at the CURDEPTH of SYSTEM.CLUSTER.COMMAND.QUEUE.

Related concepts:

“Clustering: Using REFRESH CLUSTER best practices” on page 55

You use the **REFRESH CLUSTER** command to discard all locally held information about a cluster and rebuild that information from the full repositories in the cluster. You should not need to use this command, except in exceptional circumstances. If you do need to use it, there are special considerations about how you use it. This information is a guide based on testing and feedback from customers.

Routing in publish/subscribe hierarchies

If your distributed queue manager topology is a publish/subscribe hierarchy, and a subscription is made on a queue manager, by default a proxy subscription is created on every queue manager in the hierarchy. Publications received on any queue manager are then routed through the hierarchy to each queue manager that hosts a matching subscription.

For an introduction to how messages are routed between queue managers in publish/subscribe hierarchies and clusters, see Distributed publish/subscribe networks.

When a subscription to a topic is made on a queue manager in a distributed publish/subscribe hierarchy, the queue manager manages the process by which the subscription is propagated to connected queue managers. *Proxy subscriptions* flow to all queue managers in the network. A proxy subscription gives a queue manager the information it needs to forward a publication to those queue managers that host subscriptions for that topic. Each queue manager in a publish/subscribe hierarchy is only aware of its direct relations. Publications put to one queue manager are sent, through its direct relations, to those queue managers with subscriptions. This is illustrated in the following figure, in which *Subscriber 1* registers a subscription for a particular topic on the *Asia* queue manager (1). Proxy subscriptions for this subscription on the *Asia* queue manager are forwarded to all other queue managers in the network (2,3,4).

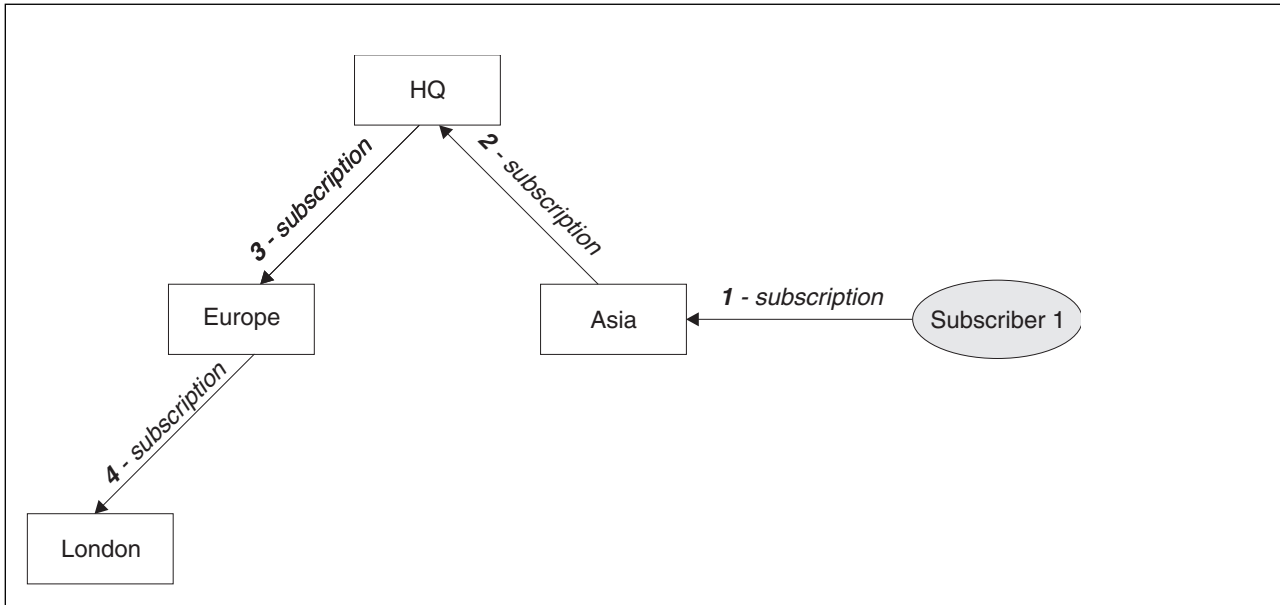


Figure 31. Propagation of subscriptions through a queue manager network

A queue manager consolidates all the subscriptions that are created on it, whether from local applications or from remote queue managers. It creates proxy subscriptions for the topics of the subscriptions with its neighbors, unless a proxy subscription already exists. This is illustrated in the following figure, in which *Subscriber 2* registers a subscription, to the same topic as in Figure 31, on the *HQ* queue manager (5). The subscription for this topic is forwarded to the *Asia* queue manager, so that it is aware that subscriptions exist elsewhere on the network (6). The subscription is not forwarded to the *Europe* queue manager, because a subscription for this topic has already been registered; see step 3 in Figure 31.

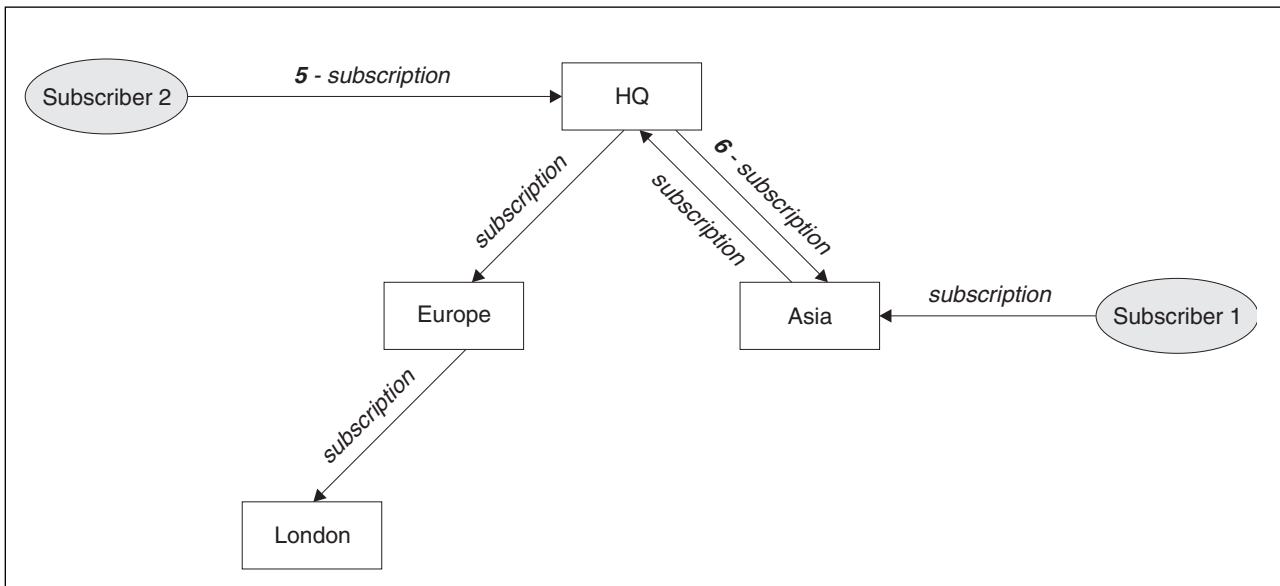


Figure 32. Multiple subscriptions

When an application publishes information to a topic, by default the receiving queue manager forwards it to all queue managers that have valid subscriptions to the topic. It might forward it through one or more intermediate queue managers. This is illustrated in the following figure, in which a publisher sends a publication, on the same topic as in Figure 32, to the *Europe* queue manager (7). A subscription for this topic exists from *HQ* to *Europe*, so the publication is forwarded to the *HQ* queue manager (8). However,

no subscription exists from *London* to *Europe* (only from *Europe* to *London*), so the publication is not forwarded to the *London* queue manager. The *HQ* queue manager sends the publication directly to *Subscriber 2* and to the *Asia* queue manager (9). The publication is forwarded to *Subscriber 1* from *Asia* (10).

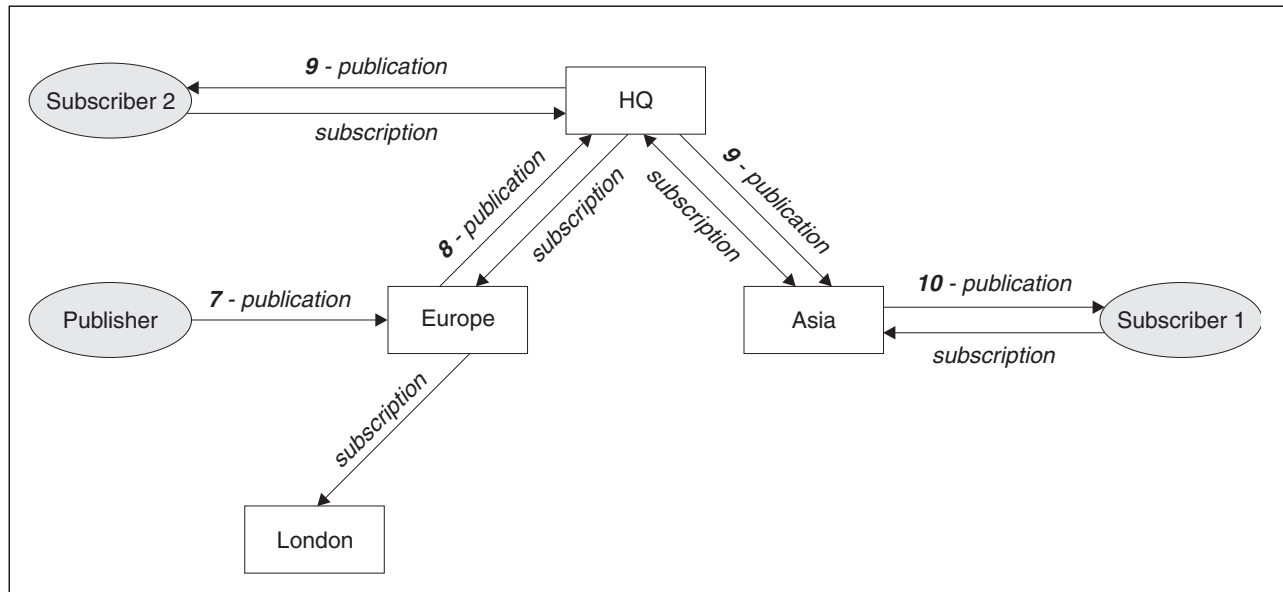


Figure 33. Propagation of publications through a queue manager network

When a queue manager sends any publications or subscriptions to another queue manager, it sets its own user ID in the message. If you are using a publish/subscribe hierarchy, and if the incoming channel is set up to put messages with the authority of the user ID in the message, then you must authorize the user ID of the sending queue manager. See Using default user IDs with a queue manager hierarchy.

Note: If you instead use publish/subscribe clusters, authorization is handled by the cluster.

Summary and additional considerations

A publish/subscribe hierarchy gives you precise control over the relationship between queue managers. After it has been created, it needs little manual intervention to administer. However it also imposes certain constraints upon your system:

- The higher nodes in the hierarchy, especially the root node, must be hosted on robust, highly available, and performant equipment. This is because more publication traffic is expected to flow through these nodes.
- The availability of every non-leaf queue manager in the hierarchy affects the ability of the network to flow messages from publishers to subscribers on other queue managers.
- By default, all topic strings subscribed to are propagated throughout the hierarchy, and publications are propagated only to remote queue managers that have a subscription to the associated topic. Therefore rapid changes to the set of subscriptions can become a limiting factor. You can change this default behavior, and instead have all publications propagated to all queue managers, which removes the need for proxy subscriptions. See Subscription performance in publish/subscribe networks.

Note: A similar restriction also applies to direct routed clusters.

- Because of the interconnected nature of publish/subscribe queue managers, it takes time for proxy subscriptions to propagate around all nodes in the network. Remote publications do not necessarily start being subscribed to immediately, so early publications might not be sent following a subscription to a new topic string. You can remove the problems caused by the subscription delay by having all

publications propagated to all queue managers, which removes the need for proxy subscriptions. See Subscription performance in publish/subscribe networks.

Note: This restriction also applies to direct routed clusters.

- For a publish/subscribe hierarchy, adding or removing queue managers requires manual configuration to the hierarchy, with careful consideration to the location of those queue managers and their reliance on other queue managers. Unless you are adding or removing queue managers that are at the bottom of the hierarchy, and therefore have no further branches below them, you will also have to configure other queue managers in the hierarchy.

Before you use a publish/subscribe hierarchy as your routing mechanism, explore the alternative approaches detailed in “Direct routing in publish/subscribe clusters” on page 63 and “Topic host routing in publish/subscribe clusters” on page 68.

Distributed publish/subscribe system queues

Four system queues are used by queue managers for publish/subscribe messaging. You need to be aware of their existence only for problem determination and capacity planning purposes.

See Balancing producers and consumers in publish/subscribe networks for guidance on how to monitor these queues.

Table 6. Publish/subscribe system queues on distributed platforms

System queue	Purpose
SYSTEM.INTER.QMGR.CONTROL	IBM MQ distributed publish/subscribe control queue
SYSTEM.INTER.QMGR.FANREQ	IBM MQ distributed publish/subscribe internal proxy subscription fan-out process input queue
SYSTEM.INTER.QMGR.PUBS	IBM MQ distributed publish/subscribe publications
SYSTEM.HIERARCHY.STATE	IBM MQ distributed publish/subscribe hierarchy relationship state

z/OS On z/OS, you set up the necessary system objects when you create the queue manager, by including the CSQ4INSR and CSQ4INSG samples in the CSQINP2 initialization input data set. For more information, see Task 13: Customize the initialization input data sets.

The attributes of the publish/subscribe system queues are shown in Table 7.

Table 7. Attributes of publish/subscribe system queues

Attribute	Default value
DEFPSIST	Yes
DEFSOPT	EXC
MAXMSGL	<p>distributed IBM i On AIX, HP-UX, Linux, IBM i, Solaris, and Windows platforms: The value of the MAXMSGL parameter of the ALTER QMGR command</p> <p>z/OS On z/OS: 104857600 (that is, 100 MB)</p>
MAXDEPTH	999999999
SHARE	N/A
z/OS STGCLASS	This attribute is used only on z/OS platforms

Note: The only queue that contains messages put by applications is SYSTEM.INTER.QMGR.PUBS. **MAXDEPTH** is set to its maximum value for this queue to allow temporary build up of published messages during outages or times of excessive load. If the queue manager is running on a system where that depth of

queue could not be contained, this should be adjusted.

Related information:

Distributed publish/subscribe troubleshooting

Distributed publish/subscribe system queue errors:

Errors can occur when distributed publish/subscribe queue manager queues are unavailable. This affects the propagation of subscription knowledge across the publish/subscribe network, and publication to subscriptions on remote queue managers.

If the fan-out request queue SYSTEM.INTER.QMGR.FANREQ is unavailable, the creation of a subscription might generate an error, and error messages will be written to the queue manager error log when proxy subscriptions need to be delivered to directly connected queue managers.

If the hierarchy relationship state queue SYSTEM.HIERARCHY.STATE is unavailable, an error message is written to the queue manager error log and the publish/subscribe engine is put into COMPAT mode. To view the publish/subscribe mode, use the command DISPLAY QMGR PSMODE.

If any other of the SYSTEM.INTER.QMGR queues are unavailable, an error message is written to the queue manager error log and, although function is not disabled, it is likely that publish/subscribe messages will build up on queues on this or remote queue managers.

If the publish/subscribe system queue or required transmission queue to a parent, child or publish/subscribe cluster queue manager is unavailable, the following outcomes occur:

- The publications are not delivered, and a publishing application might receive an error. For details of when the publishing application receives an error, see the following parameters of the **DEFINE TOPIC** command: **PMSGDLV** , **NPMSGDLV** , and **USEDLQ** .
- Received inter-queue manager publications are backed out to the input queue, and subsequently re-attempted. If the backout threshold is reached, the undelivered publications are placed on the dead letter queue. The queue manager error log will contain details of the problem.
- An undelivered proxy subscription is backed out to the fanout request queue, and subsequently attempted again. If the backout threshold is reached, the undelivered proxy subscription is not delivered to any connected queue manager, and is placed on the dead letter queue. The queue manager error log will contain details of the problem, including details of any necessary corrective administrative action required.
- Hierarchy relationship protocol messages fail, and the connection status is flagged as ERROR. To view the connection status, use the command **DISPLAY PUBSUB**.

Related information:

Distributed publish/subscribe troubleshooting

Planning your storage and performance requirements


You must set realistic and achievable storage, and performance goals for your IBM MQ system. Use the links to find out about factors that affect storage and performance on your platform.

The requirements vary depending on the systems that you are using IBM MQ on, and what components you want to use.

For the latest information about supported hardware and software environments, see the IBM MQ System Requirements Web site:

www.ibm.com/software/integration/wmq/requirements/

IBM MQ stores queue manager data in the file system. Use the following links to find out about planning and configuring directory structures for use with IBM MQ:

- “Planning file system support” on page 100
- “Requirements for shared file systems” on page 101
- “Sharing IBM MQ files” on page 110
- “Directory structure on UNIX and Linux systems” on page 113
- “Directory structure on Windows systems” on page 122
-  “Directory structure on IBM i” on page 125

Use the following links for information about system resources, shared memory, and process priority on UNIX and Linux:

- “IBM MQ and UNIX System V IPC resources” on page 128
- “Shared memory on AIX” on page 128
- “IBM MQ and UNIX Process Priority” on page 129


Use the following link for information about log files:

- Calculating the size of the log

Related concepts:

“Planning” on page 1

When planning your IBM MQ environment, consider the support that IBM MQ provides for single and multiple queue manager architectures, and for point-to-point and publish/subscribe messaging styles. Also plan your resource requirements, and your use of logging and backup facilities.

 “Planning your IBM MQ environment on z/OS” on page 132

When planning your IBM MQ environment, you must consider the resource requirements for data sets, page sets, Db2, Coupling Facilities, and the need for logging, and backup facilities. Use this topic to plan the environment where IBM MQ runs.

Related information:

Hardware and software requirements on UNIX and Linux

Hardware and software requirements on Windows

Disk space requirements on distributed platforms



The storage requirements for IBM MQ depend on which components you install, and how much working space you need.


Disk storage is required for the optional components you choose to install, including any prerequisite components they require. The total storage requirement also depends on the number of queues that you use, the number and size of the messages on the queues, and whether the messages are persistent. You also require archiving capacity on disk, tape, or other media, as well as space for your own application programs.

The following table shows the approximate disk space required when you install various combinations of the product on different platforms. (Values are rounded up to the nearest 5 MB, where a MB is 1,048,576 bytes.)





Notes:

1. On IBM i you cannot separate the native client from the server. The server figure in the table is for 5724H72*BASE without Java™, together with the English Language Load (2924). There are 22 possible unique language loads.
2. The figure in the table is for the native client 5725A49 *BASE without Java.
3. Java and JMS classes can be added to both server and client bindings. If you want to include these features add 110 MB.
4. Adding samples source to either the client or server adds an extra 10 MB.
5. Adding samples to Java and JMS classes adds an extra 5 MB.

Platform	Client installation ¹	Server installation ²	IBM MQ MFT installation ³	Full installation ⁴
AIX	145 MB	190 MB	705 MB	915 MB
HP-UX	225 MB	310 MB	1075 MB	1340 MB
IBM i	215 MB	450 MB	80 MB	655 MB
Linux for System x (32 bit)	85 MB	N/A	N/A	120 MB
Linux for System x (64 bit)	125 MB	170 MB	575 MB	935 MB
 Linux on POWER® Systems - Little Endian	90 MB	125 MB	555 MB	620 MB
Linux on POWER Systems - Big Endian	130 MB	170 MB	565 MB	715 MB
Linux for IBM Z	125 MB	160 MB	560 MB	665 MB
Solaris x86-64, AMD64, EM64T, and compatible processors	105 MB ⁵	150 MB ⁵	695 MB	860 MB
Solaris SPARC	105 MB ⁵	150 MB ⁵	680 MB	820 MB
Windows (32 bit install) ⁶	390 MB	N/A	N/A	475 MB
Windows (64 bit install) ⁶	445 MB	555 MB	710 MB	1005 MB

Usage notes

1. A client installation includes the following components:
 - Runtime
 - Client

2. A server installation includes the following components:
 - Runtime
 - Server
3. An IBM MQ Managed File Transfer installation includes the following components:
 - IBM MQ Managed File Transfer Service, Logger, Agent, Tools, and Base components
 - Runtime
 - Server
 - Java
 - JRE
4. A full installation includes all available components.
5.  On Solaris platforms you must install silently to get this combination of components.
6.  Not all the components listed here are installable features on Windows systems; their functionality is sometimes included in other features. See IBM MQ features for Windows systems.

Related information:

Choosing what to install

Planning file system support


Queue manager data is stored in the file system. A queue manager makes use of file system locking to prevent multiple instances of a multi-instance queue manager being active at the same time.

Shared file systems

Shared file systems enable multiple systems to access the same physical storage device concurrently. Corruption would occur if multiple systems accessed the same physical storage device directly without some means of enforcing locking and concurrency control. Operating systems provide local file systems with locking and concurrency control for local processes; network file systems provide locking and concurrency control for distributed systems.

Historically, networked file systems have not performed fast enough, or provided sufficient locking and concurrency control, to meet the requirements for logging messages. Today, networked file systems can provide good performance, and implementations of reliable network file system protocols such as *RFC 3530, Network File System (NFS) version 4 protocol*, meet the requirements for logging messages reliably.

Shared file systems and IBM MQ

Queue manager data for a multi-instance queue manager is stored in a shared network file system. On Microsoft Windows, UNIX and Linux systems, the queue manager's data files and log files must be placed in shared network file system.  On IBM i, journals are used instead of log files, and journals cannot be shared. Multi-instance queue managers on IBM i use journal replication, or switchable journals, to make journals available between different queue manager instances.



Prior to release v7.0.1, IBM MQ does not support queue manager data stored on networked storage accessed as a shared file system. If queue manager data is placed on shared networked storage, then you need to ensure the queue manager data is not accessed by another instance of the queue manager running at the same time.

From v7.0.1 onwards, IBM MQ uses locking to prevent multiple instances of the same multi-instance queue manager being active at the same time. The same locking also ensures that two separate queue managers cannot inadvertently use the same set of queue manager data files. Only one instance of a queue manager can have its lock at a time. Consequently, IBM MQ does support queue manager data stored on networked storage accessed as a shared file system.

Because not all the locking protocols of network file systems are robust, and because a file system might be configured for performance rather than data integrity, you must run the **amqmfscck** command to test whether a network file system will control access to queue manager data and logs correctly. This command applies only UNIX and IBM i systems. On Microsoft Windows, there is only one supported network file system and the **amqmfscck** command is not required.

Related tasks:

“Verifying shared file system behavior” on page 102

Run **amqmfscck** to check whether a shared file system on UNIX  **IBM i**  and IBM i systems meets the requirements for storing the queue manager data of a multi-instance queue manager. Run the IBM MQ MQI client sample program **amqsfhac** in parallel with **amqmfscck** to demonstrate that a queue manager maintains message integrity during a failure.

Requirements for shared file systems


Shared file systems must provide data write integrity, guaranteed exclusive access to files and release locks on failure to work reliably with IBM MQ.

There are three fundamental requirements that a shared file system must meet to log messages reliably:

1. Data write integrity.

Data write integrity is sometimes called "Write through to disk on flush". The queue manager must be able to synchronize with data being successfully committed to the physical device. In a transactional system, you need to be sure that some writes have been safely committed before continuing with other processing.

More specifically, IBM MQ on UNIX platforms uses the `O_SYNC` open option and the `fsync()` system call to explicitly force writes to recoverable media, and the write operation is dependent upon these options operating correctly.

Attention:  You should mount the file system with the `async` option, which still supports the option of synchronous writes and gives better performance than the `sync` option.

Note, however, that if the file system has been exported from Linux, you must still export the file system using the `sync` option.

2. Guaranteed exclusive access to files.

In order to synchronize multiple queue managers, there needs to be a mechanism for a queue manager to obtain an exclusive lock on a file.

3. Release locks on failure.

If a queue manager fails, or if there is a communication failure with the file system, files locked by the queue manager need to be unlocked and made available to other processes without waiting for the queue manager to be reconnected to the file system.

For multi-instance queue managers on Microsoft Windows, the networked storage must be accessed by the Common Internet File System (CIFS) protocol used by Microsoft Windows networks.

For multi-instance queue managers on other supported platforms, the storage must be accessed by a network file system protocol which is Posix-compliant and supports lease-based locking. Network File System Version 4 satisfies this requirement. Older file systems, such as Network File System Version 3, which do not have a reliable mechanism to release locks after a failure, must not be used with multi-instance queue managers.

A shared file system *must* meet these requirements for IBM MQ to operate reliably. If it does not, the queue manager data and logs get corrupted when using the shared file system in a multi-instance queue manager configuration.

You must check whether the shared file system you plan to use meets these requirements. You must also check whether the file system is correctly configured for reliability. Shared file systems sometimes provide configuration options to improve performance at the expense of reliability.

For further information, see [Testing and support statement for IBM MQ multi-instance queue managers](#).

Under normal circumstances IBM MQ operates correctly with attribute caching and it is not necessary to disable caching, for example by setting NOAC on an NFS mount. Attribute caching can cause issues when multiple file system clients are contending for write access to the same file on the file system server, as the cached attributes used by each client might not be the same as those attributes on the server. An example of files accessed in this way are queue manager error logs for a multi-instance queue manager. The queue manager error logs might be written to by both an active and a standby queue manager instance and cached file attributes might cause the error logs to grow larger than expected, before rollover of the files occurs.

To help to check the file system, run the task [Verifying shared file system behavior](#). This task checks if your shared file system meets requirements 2 and 3. You need to verify requirement 1 in your shared file system documentation, or by experimenting with logging data to the disk.

Disk faults can cause errors when writing to disk, which IBM MQ reports as First Failure Data Capture errors. You can run the file system checker for your operating system to check the shared file system for any disk faults. For example, on UNIX systems the file system checker is called `fsck`. On Windows platforms the file system checker is called `CHKDSK`, or `SCANDISK`.

Note: You should put only queue manager data on an NFS server. On the NFS, use the following three options with the mount command to make the system secure:

noexec


By using this option, you stop binary files from being run on the NFS, which prevents a remote user from running unwanted code on the system.

nosuid

By using this option, you prevent the use of the set-user-identifier and set-group-identifier bits, which prevents a remote user from gaining higher privileges.

nodev By using this option, you stop character and block special devices from being used or defined, which prevents a remote user from getting out of a chroot jail.

Verifying shared file system behavior:

Run `amqmfsc` to check whether a shared file system on UNIX  and IBM i systems meets the requirements for storing the queue manager data of a multi-instance queue manager. Run the IBM MQ MQI client sample program `amqsfhac` in parallel with `amqmfsc` to demonstrate that a queue manager maintains message integrity during a failure.

Before you begin

You need a server with networked storage, and two other servers connected to it that have IBM MQ installed. You must have administrator (root) authority to configure the file system, and be an IBM MQ Administrator to run `amqmfsc`.

About this task

“Requirements for shared file systems” on page 101 describes the file system requirements for using a shared file system with multi-instance queue managers. The IBM MQ technote [Testing and support statement for IBM MQ multi-instance queue managers](#) lists the shared file systems that IBM has already tested with. The procedure in this task describes how to test a file system to help you assess whether an unlisted file system maintains data integrity.

Failover of a multi-instance queue manager can be triggered by hardware or software failures, including networking problems which prevent the queue manager writing to its data or log files. Mainly, you are

interested in causing failures on the file server. But you must also cause the IBM MQ servers to fail, to test any locks are successfully released. To be confident in a shared file system, test all of the following failures, and any other failures that are specific to your environment:

1. Shutting down the operating system on the file server including syncing the disks.
2. Halting the operating system on the file server without syncing the disks.
3. Pressing the reset button on each of the servers.
4. Pulling the network cable out of each of the servers.
5. Pulling the power cable out of each of the servers.
6. Switching off each of the servers.

Create the directory on the networked storage that you are going to use to share queue manager data and logs. The directory owner must be an IBM MQ Administrator, or in other words, a member of the `mqm` group on UNIX. The user who runs the tests must have IBM MQ Administrator authority.

Use the example of exporting and mounting a file system in [Create a multi-instance queue manager on Linux](#) or [Mirrored journal configuration on an ASP using ADDMQMJRN](#) to help you through configuring the file system. Different file systems require different configuration steps. Read the file system documentation.

Procedure

In each of the checks, cause all the failures in the previous list while the file system checker is running. If you intend to run `amqsfhac` at the same time as `amqmfscck`, do the task, “Running `amqsfhac` to test message integrity” on page 108 in parallel with this task.

1. Mount the exported directory on the two IBM MQ servers.

On the file system server create a shared directory `shared`, and a subdirectory to save the data for multi-instance queue managers, `qmdata`. For an example of setting up a shared directory for multi-instance queue managers on Linux, see [Example in Create a multi-instance queue manager on Linux](#)

2. Check basic file system behavior.

On one IBM MQ server, run the file system checker with no parameters.

```
amqmfscck /shared/qmdata
```

Figure 34. On IBM MQ server 1

3. Check concurrently writing to the same directory from both IBM MQ servers.

On both IBM MQ servers, run the file system checker at the same time with the `-c` option.

```
amqmfscck -c /shared/qmdata
```

Figure 35. On IBM MQ server 1

```
amqmfscck -c /shared/qmdata
```

Figure 36. On IBM MQ server 2

4. Check waiting for and releasing locks on both IBM MQ servers.

On both IBM MQ servers run the file system checker at the same time with the `-w` option.

```
amqmfscck -w /shared/qmdata
```

Figure 37. On IBM MQ server 1

```
amqmfscck -w /shared/qmdata
```

Figure 38. On IBM MQ server 2

5. Check for data integrity.

a. Format the test file.

Create a large file in the directory being tested. The file is formatted so that the subsequent phases can complete successfully. The file must be large enough that there is sufficient time to interrupt the second phase to simulate the failover. Try the default value of 262144 pages (1 GB). The program automatically reduces this default on slow file systems so that formatting completes in about 60 seconds

```
amqmfscck -f /shared/qmdata
```

The server responds with the following messages:
Formatting test file for data integrity test.

Test file formatted with 262144 pages of data.

Figure 39. On IBM MQ server 1

b. Write data into the test file using the file system checker while causing a failure.

Run the test program on two servers at the same time. Start the test program on the server which is going to experience the failure, then start the test program on the server that is going to survive the failure. Cause the failure you are investigating.

The first test program stops with an error message. The second test program obtains the lock on the test file and writes data into the test file starting where the first test program left off. Let the second test program run to completion.

Table 8. Running the data integrity check on two servers at the same time

IBM MQ server 1	IBM MQ server 2
amqmfscck -a /shared/qmdata	

Table 8. Running the data integrity check on two servers at the same time (continued)

IBM MQ server 1	IBM MQ server 2
Please start this program on a second machine with the same parameters.	amqmfscck -a /shared/qmdata
File lock acquired.	Waiting for lock...
Start a second copy of this program with the same parameters on another server.	Waiting for lock...
Writing data into test file.	Waiting for lock...
To increase the effectiveness of the test, interrupt the writing by ending the process, temporarily breaking the network connection to the networked storage, rebooting the server or turning off the power.	Waiting for lock...
Turn the power off here.	Waiting for lock...
	File lock acquired.
	Reading test file
	Checking the integrity of the data read.
	Appending data into the test file after data already found.
	The test file is full of data. It is ready to be inspected for data integrity.

The timing of the test depends on the behavior of the file system. For example, it typically takes 30 - 90 seconds for a file system to release the file locks obtained by the first program following a power outage. If you have too little time to introduce the failure before the first test program has filled the file, use the -x option of **amqmfscck** to delete the test file. Try the test from the start with a larger test file.

- c. Verify the integrity of the data in the test file.

```
amqmfscck -i /shared/qmdata
```

The server responds with the following messages:

```
File lock acquired
```

```
Reading test file checking the integrity of the data read.
```

```
The data read was consistent.
```

```
The tests on the directory completed successfully.
```

Figure 40. On IBM MQ server 2

6. Delete the test files.

```
amqmfscck -x /shared/qmdata
```

```
Test files deleted.
```

Figure 41. On IBM MQ server 2

The server responds with the message:

```
Test files deleted.
```

Results

The program returns an exit code of zero if the tests complete successfully, and non-zero otherwise.

Examples

The first set of three examples shows the command producing minimal output.

Successful test of basic file locking on one server

```
> amqmfscck /shared/qmdata
```

```
The tests on the directory completed successfully.
```

Failed test of basic file locking on one server

```
> amqmfscck /shared/qmdata
```

```
AMQ6245: Error Calling 'write()[2]' on file '/shared/qmdata/amqmfscck.lck' error '2'.
```

Successful test of locking on two servers

Table 9. Successful locking on two servers

IBM MQ server 1	IBM MQ server 2
<pre>> amqmfscck -w /shared/qmdata Please start this program on a second machine with the same parameters. Lock acquired. Press Return or terminate the program to release the lock.</pre>	
	<pre>> amqmfscck -w /shared/qmdata Waiting for lock...</pre>
<pre>[Return pressed] Lock released.</pre>	

Table 9. Successful locking on two servers (continued)

IBM MQ server 1	IBM MQ server 2
	Lock acquired. The tests on the directory completed successfully

The second set of three examples shows the same commands using verbose mode.

Successful test of basic file locking on one server

```
> amqmfscck -v /shared/qmdata
System call: stat("/shared/qmdata")
System call: fd = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fchmod(fd, 0666)
System call: fstat(fd)
System call: fcntl(fd, F_SETLK, F_WRLCK)
System call: write(fd)
System call: close(fd)
System call: fd = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fd, F_SETLK, F_WRLCK)
System call: close(fd)
System call: fd1 = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fd1, F_SETLK, F_RDLCK)
System call: fd2 = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fd2, F_SETLK, F_RDLCK)
System call: close(fd2)
System call: write(fd1)
System call: close(fd1)
The tests on the directory completed successfully.
```

Failed test of basic file locking on one server

```
> amqmfscck -v /shared/qmdata
System call: stat("/shared/qmdata")
System call: fd = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fchmod(fd, 0666)
System call: fstat(fd)
System call: fcntl(fd, F_SETLK, F_WRLCK)
System call: write(fd)
System call: close(fd)
System call: fd = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fd, F_SETLK, F_WRLCK)
System call: close(fd)
System call: fd = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fd, F_SETLK, F_RDLCK)
System call: fdSameFile = open("/shared/qmdata/amqmfscck.lck", O_RDWR, 0666)
System call: fcntl(fdSameFile, F_SETLK, F_RDLCK)
System call: close(fdSameFile)
System call: write(fd)
AMQxxxx: Error calling 'write()[2]' on file '/shared/qmdata/amqmfscck.lck', errno 2
(Permission denied).
```

Successful test of locking on two servers

Table 10. Successful locking on two servers - verbose mode

IBM MQ server 1	IBM MQ server 2
<pre>> amqmfscck -wv /shared/qmdata Calling 'stat("/shared/qmdata")' Calling 'fd = open("/shared/qmdata/amqmfscck.1kw", O_EXCL O_CREAT O_RDWR, 0666)' Calling 'fchmod(fd, 0666)' Calling 'fstat(fd)' Please start this program on a second machine with the same parameters. Calling 'fcntl(fd, F_SETLK, F_WRLCK)' Lock acquired. Press Return or terminate the program to release the lock.</pre>	
	<pre>> amqmfscck -wv /shared/qmdata Calling 'stat("/shared/qmdata")' Calling 'fd = open("/shared/qmdata/amqmfscck.1kw", O_EXCL O_CREAT O_RDWR,0666)' Calling 'fd = open("/shared/qmdata/amqmfscck.1kw", O_RDWR, 0666)' Calling 'fcntl(fd, F_SETLK, F_WRLCK) 'Waiting for lock...</pre>
<pre>[Return pressed] Calling 'close(fd)' Lock released.</pre>	
	<pre>Calling 'fcntl(fd, F_SETLK, F_WRLCK)' Lock acquired. The tests on the directory completed successfully</pre>

Related information:

amqmfscck (file system check)

Running amqsfhac to test message integrity:

amqsfhac checks that a queue manager using networked storage maintains data integrity following a failure.

Before you begin

You require four servers for this test. Two servers for the multi-instance queue manager, one for the file system, and one for running **amqsfhac** as a IBM MQ MQI client application.

Follow step 1 on page 103 in Procedure to set up the file system for a multi-instance queue manager.

About this task

Procedure

1. Create a multi-instance queue manager on another server, QM1, using the file system you created in step 1 on page 103 in Procedure.
 - See Create a multi-instance queue manager.
2. Start the queue manager on both servers making it highly available.
 - On server 1:

```
strmqm -x QM1
```
 - On server 2:

```
strmqm -x QM1
```
3. Set up the client connection to run **amqsfhac**.

- a. Use the procedure in Verifying a client installation to set up a client connection, or the example scripts in Reconnectable client samples.
- b. Modify the client channel to have two IP addresses, corresponding to the two servers running QM1.

In the example script, modify:

```
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNAME('LOCALHOST(2345)') QMNAME(QM1) REPLACE
```

To:

```
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNAME('server1(2345),server2(2345)') QMNAME(QM1) REPLACE
```

Where server1 and server2 are the host names of the two servers, and 2345 is the port that the channel listener is listening on. Usually this defaults to 1414. You can use 1414 with the default listener configuration.

4. Create two local queues on QM1 for the test. Run the following MQSC script:

```
DEFINE QLOCAL(TARGETQ) REPLACE
DEFINE QLOCAL(SIDEQ) REPLACE
```

5. Test the configuration with **amqsfhac**

```
amqsfhac QM1 TARGETQ SIDEQ 2 2 2
```

6. Test message integrity while you are testing file system integrity.

Run **amqsfhac** during step 5 on page 104 of Procedure.

```
amqsfhac QM1 TARGETQ SIDEQ 10 20 0
```

If you stop the active queue manager instance, **amqsfhac** reconnects to the other queue manager instance once it has become active. Restart the stopped queue manager instance again, so that you can reverse the failure in your next test. You will probably need to increase the number of iterations based on experimentation with your environment so that the test program runs for sufficient time for the failover to occur.

Results

An example of running **amqsfhac** in step 6 is shown in Figure 42 on page 110. The test is a success.

If the test detected a problem, the output would report the failure. In some test runs MQRC_CALL_INTERRUPTED might report "Resolving to backed out". It makes no difference to the result. The outcome depends on whether the write to disk was committed by the networked file storage before or after the failure took place.

```
Sample AMQSFHAC start
qmname = QM1
qname = TARGETQ
sidename = SIDEQ
transize = 10
iterations = 20
verbose = 0
Iteration 0
Iteration 1
Iteration 2
Iteration 3
Iteration 4
Iteration 5
Iteration 6
Resolving MQRC_CALL_INTERRUPTED
MQGET browse side tranid=14 pSideinfo->tranid=14
Resolving to committed
Iteration 7
Iteration 8
Iteration 9
Iteration 10
Iteration 11
Iteration 12
Iteration 13
Iteration 14
Iteration 15
Iteration 16
Iteration 17
Iteration 18
Iteration 19
Sample AMQSFHAC end
```

Figure 42. Output from a successful run of amqsfhac

Related information:

High availability sample programs

Sharing IBM MQ files

Some IBM MQ files are accessed exclusively by an active queue manager, other files are shared.

IBM MQ files are split into program files and data files. Program files are typically installed locally on each server running IBM MQ. Queue managers share access to data files and directories in the default data directory. They require exclusive access to their own queue manager directory trees contained in each of the qmgrs and log directories shown in Figure 43 on page 111.

Figure 43 on page 111 is a high-level view of the IBM MQ directory structure. It shows the directories which can be shared between queue managers and made remote. The details vary by platform. The dotted lines indicate configurable paths.




Figure 43. Overall view of IBM MQ directory structure

Program files

The program files directory is typically left in the default location, is local, and shared by all the queue managers on the server.

Data files

The data files directory is typically local in the default location, `/var/mqm` on UNIX and Linux systems and configurable on installation on Windows. It is shared between queue managers. You can make the default location remote, but do not share it between different installations of IBM MQ. The `DefaultPrefix` attribute in the IBM MQ configuration points to this path.

qmgrs From v7.0.1, there are two alternative ways to specify the location of queue manager data.

Using Prefix

The `Prefix` attribute specifies the location of the `qmgrs` directory. IBM MQ constructs the queue manager directory name from the queue manager name and creates it as a subdirectory of the `qmgrs` directory.

The `Prefix` attribute is located in the `QueueManager` stanza, and is inherited from the value in the `DefaultPrefix` attribute. By default, for administrative simplicity, queue managers typically share the same `qmgrs` directory.

The `QueueManager` stanza is in the `mqs.ini` file.

If you change the location of the `qmgrs` directory for any queue manager, you must change the value of its `Prefix` attribute.

The `Prefix` attribute for the `QM1` directory in Figure 43 for a UNIX and Linux platform is, `Prefix=/var/mqm`

Using DataPath

The `DataPath` attribute specifies the location of the queue manager data directory.

The `DataPath` attribute specifies the complete path, including the name of the queue manager data directory. The `DataPath` attribute is unlike the `Prefix` attribute, which specifies an incomplete path to the queue manager data directory.

The `DataPath` attribute, if it is specified, is located in the `QueueManager` stanza. If it has been specified, it takes precedence over any value in the `Prefix` attribute.

The `QueueManager` stanza is in the `mqs.ini` file.

If you change the location of the queue manager data directory for any queue manager you must change the value of the `DataPath` attribute.

The `DataPath` attribute for the QM1 directory in Figure 43 on page 111 for a UNIX or Linux platform is,

```
DataPath=/var/mqm/qmgrs/QM1
```

Log

The log directory is specified separately for each queue manager in the Log stanza in the queue manager configuration. The queue manager configuration is in `qm.ini`.

DataPath/QmgrName/@IPCC subdirectories

The *DataPath/QmgrName/@IPCC* subdirectories are in the shared directory path. They are used to construct the directory path for IPC file system objects. They need to distinguish the namespace of a queue manager when a queue manager is shared between systems. Before V7.0.1, a queue manager was only used on one system. One set of subdirectories was sufficient to define the directory path to IPC file system objects, see Figure 44.

```
DataPath/QmgrName/@IPCC/esem
```

Figure 44. Example IPC subdirectory, pre-V7.0.1

In V7.0.1, and later, the IPC file system objects have to be distinguished by system. A subdirectory, for each system the queue manager runs on, is added to the directory path, see Figure 45.

```
DataPath/QmgrName/@IPCC/esem/myHostName/
```

Figure 45. Example IPC subdirectory, V7.0.1 and subsequent releases

myHostName is up to the first 20 characters of the host name returned by the operating system. On some systems, the host name might be up to 64 characters in length before truncation. The generated value of *myHostName* might cause a problem for two reasons:

1. The first 20 characters are not unique.
2. The host name is generated by a DHCP algorithm that does not always allocate the same host name to a system.

In these cases, set *myHostName* using the environment variable, `MQC_IPC_HOST` ; see Figure 46.

```
export MQC_IPC_HOST= myHostName
```

Figure 46. Example: setting `MQC_IPC_HOST`

Other files and directories

Other files and directories, such as the directory containing trace files, and the common error log, are normally shared and kept on the local file system.

Up until v7.0.1, IBM MQ relied upon external management to guarantee queue managers exclusive access to the queue manager data and log files. From v7.0.1 onwards, with support of shared file systems, IBM MQ manages exclusive access to these files using file system locks. A file system lock allows only one instance of a particular queue manager to be active at a time.

When you start the first instance of a particular queue manager it takes ownership of its queue manager directory. If you start a second instance, it can only take ownership if the first instance has stopped. If the first queue manager is still running, the second instance fails to start, and reports the queue manager is running elsewhere. If the first queue manager has stopped, then the second queue manager takes over ownership of the queue manager files and becomes the running queue manager.

You can automate the procedure of the second queue manager taking over from the first. Start the first queue manager with the `strmqm -x` option that permits another queue manager to take over from it. The

second queue manager then waits until the queue manager files are unlocked before attempting to take over ownership of the queue manager files, and start.

Directory structure on UNIX and Linux systems

The IBM MQ directory structure on UNIX and Linux systems can be mapped to different file systems for easier management, better performance, and better reliability.

Use the flexible directory structure of IBM MQ to take advantage of shared file systems for running multi-instance queue managers.

Use the command `crtmqm QM1` to create the directory structure shown in Figure 47 where R is the release of the product. It is a typical directory structure for a queue manager created on an IBM MQ system from IBM WebSphere MQ Version 7.0.1 onwards. Some directories, files and .ini attribute settings are omitted for clarity, and another queue manager name might be altered by mangling. The names of the file systems vary on different systems.

In a typical installation, every queue manager that you create points to common log and qmgrs directories on the local file system. In a multi-instance configuration, the log and qmgrs directories are on a network file system shared with another installation of IBM MQ.

Figure 47 shows the default configuration for IBM MQ v7.R on AIX where R is the release of the product. For examples of alternative multi-instance configurations, see “Example directory configurations on UNIX and Linux systems” on page 117.




Figure 47. Example default IBM MQ directory structure for UNIX and Linux systems

The product is installed into `/usr/mqm` on AIX and `/opt/mqm` on the other systems, by default. The working directories are installed into the `/var/mqm` directory.

Note: If you created the `/var/mqm` file system prior to installing IBM MQ, ensure that the `mqm` user has full directory permissions, for example, file mode 755.

Note: The `/var/mqm/errors` directory should be a separate filesystem to prevent FFDCs produced by the queue manager from filling the filesystem that contains `/var/mqm`.

See Creating file systems on UNIX and Linux systems for more information.

The log and qmgrs directories are shown in their default locations as defined by the default values of the LogDefaultPath and DefaultPrefix attributes in the mqs.ini file. When a queue manager is created, by default the queue manager data directory is created in *DefaultPrefix/qmgrs*, and the log file directory in *LogDefaultPath/log*. LogDefaultPath and DefaultPrefix only effects where queue managers and log files are created by default. The actual location of a queue manager directory is saved in the mqs.ini file, and the location of the log file directory is saved in the qm.ini file.

The log file directory for a queue manager is defined in the qm.ini file in the LogPath attribute. Use the -ld option on the **crtmqm** command to set the LogPath attribute for a queue manager; for example, **crtmqm -ld LogPath QM1** . If you omit the ld parameter the value of LogDefaultPath is used instead.

The queue manager data directory is defined in the DataPath attribute in the QueueManager stanza in the mqs.ini file. Use the -md option on the **crtmqm** command to set the DataPath for a queue manager; for example, **crtmqm -md DataPath QM1** . If you omit the md parameter the value of the DefaultPrefix or Prefix attribute is used instead. Prefix takes precedence over DefaultPrefix.

Typically, create QM1 specifying both the log and data directories in a single command.

```
crtmqm  
-md DataPath -ld  
LogPath QM1
```

You can modify the location of a queue manager log and data directories of an existing queue manager by editing the DataPath and LogPath attributes in the qm.ini file when the queue manager is stopped.

The path to the errors directory, like the paths to all the other directories in /var/mqm, is not modifiable. However the directories can be mounted on different file systems, or symbolically linked to different directories.

Directory content on UNIX and Linux systems:

Content of the directories associated with a queue manager.

For information about the location of the product files, see Choosing an installation location

For information about the location of the product files, see Choosing an installation location

For information about alternative directory configurations, see “Planning file system support” on page 100.

In Figure 48 on page 115, the layout is representative of IBM MQ after a queue manager has been in use for some time. The actual structure that you have depends on which operations have occurred on the queue manager.

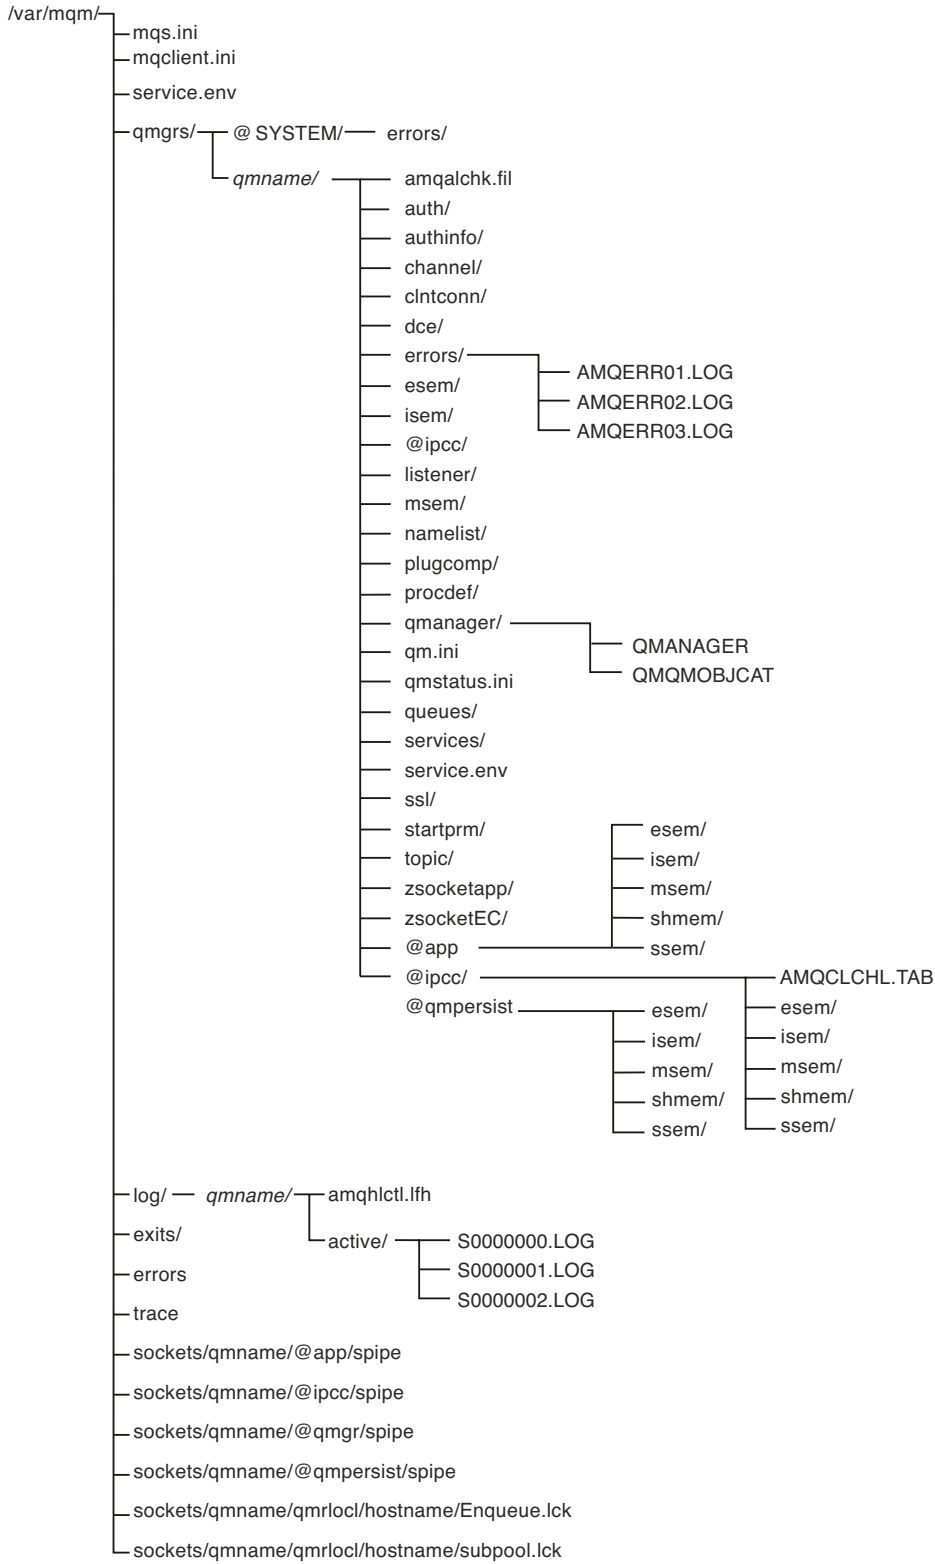


Figure 48. Default directory structure (UNIX systems) after a queue manager has been started

/var/mqm/

The /var/mqm directory contains configuration files and output directories that apply to an IBM MQ installation as a whole, and not to an individual queue manager.

Table 11. Documented content of the /var/mqm directory on UNIX systems

mqs.ini	IBM MQ installation wide configuration file, read when a queue manager starts. File path modifiable using the AMQ_MQS_INI_LOCATION environment variable. Ensure this is set and exported in the shell in which the strmqm command is run.
mqclient.ini	Default client configuration file read by IBM MQ MQI client programs. File path modifiable using the MQCLNTCF environment variable.
service.env	Contains machine scope environment variables for a service process. File path fixed.
errors/	Machine scope error logs, and FFST™ files. Directory path fixed. See also, FFST: IBM MQ for UNIX and Linux systems.
sockets/	Contains information for each queue manager for system use only.
trace/	Trace files. Directory path fixed.
exits/	Default directory containing user channel exit programs. Location modifiable in ApiExit stanzas in the mqs.ini file.
exits64/	

/var/mqm/qmgrs/qmname/

/var/mqm/qmgrs/qmname/ contains directories and files for a queue manager. The directory is locked for exclusive access by the active queue manager instance. The directory path is directly modifiable in the mqs.ini file, or by using the **md** option of the **crtmqm** command.

Table 12. Documented contents of the /var/mqm/qmgrs/qmname directory on UNIX systems

qm.ini	Queue manager configuration file, read when a queue manager starts.
errors/	Queue manager scope error logs. <i>qmname</i> = @system contains channel-related messages for an unknown or unavailable queue manager.
@ipcc/ AMQCLCHL.TAB	Default client channel control table, created by the IBM MQ server, and read by IBM MQ MQI client programs. File path modifiable using the MQCHLLIB and MQCHLTAB environment variables.
qmanager	Queue manager object file: QMANAGER Queue manager object catalog: QMQMOBJCAT

Table 12. Documented contents of the `/var/mqm/qmgrs/qmname` directory on UNIX systems (continued)

authinfo/	<p>Each object defined within the queue manager is associated with a file in these directories.</p> <p>The file name approximately matches the definition name; see, Understanding IBM MQ file names.</p>
channel/	
clntconn/	
listener/	
namelist/	
procdef/	
queues/	
services/	
topics/	
...	Other directories used by IBM MQ, such as @ipcc, to be modified only by IBM MQ.

`/var/mqm/log/qmname/`

`/var/mqm/log/qmname/` contains the queue manager log files. The directory is locked for exclusive access by the active queue manager instance. The directory path is modifiable in the `qm.ini` file, or by using the `ld` option of the `crtmqm` command.

Table 13. Documented contents of the `/var/mqm/log/qmname` directory on UNIX systems

amqhlctl.lfh	Log control file.
active/	This directory contains the log files numbered S0000000.LOG, S0000001.LOG, S0000002.LOG, and so on.

`opt/mqm`

`opt/mqm` is, by default, the installation directory on most platforms. See “Disk space requirements on distributed platforms” on page 98 for more information on the amount of space that you need for the installation directory on the platform, or platforms, that your enterprise uses.

Example directory configurations on UNIX and Linux systems:

Examples of alternative file system configurations on UNIX and Linux systems.

You can customize the IBM MQ directory structure in various ways to achieve a number of different objectives.

- Place the `qmgrs` and `log` directories on remote shared file systems to configure a multi-instance queue manager.
- Use separate file systems for the data and log directories, and allocate the directories to different disks, to improve performance by reducing I/O contention.
- Use faster storage devices for directories that have a greater effect on performance. Physical device latency is frequently a more important factor in the performance of persistent messaging than whether a device is mounted locally or remotely. The following list shows which directories are most and least performance sensitive.
 1. `log`
 2. `qmgrs`
 3. Other directories, including `/usr/mqm`
- Create the `qmgrs` and `log` directories on file systems that are allocated to storage with good resilience, such as a redundant disk array, for example.

- It is better to store the common error logs in `var/mqm/errors`, locally, rather than on a network file system, so that error relating to the network file system can be logged.

Figure 49 is a template from which alternative IBM MQ directory structures are derived. In the template, dotted lines represent paths that are configurable. In the examples, the dotted lines are replaced by solid lines that correspond to the configuration information stored in the `AMQ_MQS_INI_LOCATION` environment variable, and in the `mqs.ini` and `qm.ini` files.

Note: The path information is shown as it appears in the `mqs.ini` or `qm.ini` files. If you supply path parameters in the `crtmqm` command, omit the name of the queue manager directory: the queue manager name is added to the path by IBM MQ after it has been mangled.




Figure 49. Directory structure pattern template

Examples of configured directory structures follow. The first example shows a typical default directory structure for IBM MQ v7.0.1 created by issuing the `crtmqm QM1` command. The second example shows how a typical directory structure appears for a queue manager created using an IBM MQ release earlier than v7.0.1. The directory structure does not change.

Queue managers newly created in Version 7.0.1 have a different configuration file to earlier releases of v7. If you need to remove the v7.0.1 fix pack to revert to v7.0.0.2, you need to re-create the configuration files. You might only need to use the `Prefix` attribute to define the path to the new queue manager data directory, or you might need to move the queue manager data directory and log directories to a different location. The safest way to reconfigure the queue manager is to save the queue manager data and log directories, delete and re-create the queue manager, and then replace the data and log directories in their new location, with the ones that have been saved.

Typical directory structure for release v7.0.1 onwards

Figure 50 on page 119 is the default directory structure created in v7.0.1 by issuing the command `crtmqm QM1`.

The `mqs.ini` file has a stanza for the QM1 queue manager, created by referring to the value of `DefaultPrefix`. The `Log` stanza in the `qm.ini` file has a value for `LogPath`, set by reference to `LogDefaultPath` in `mqs.ini`.

Use the optional `crtmqm` parameters to override the default values of `DataPath` and `LogPath`.

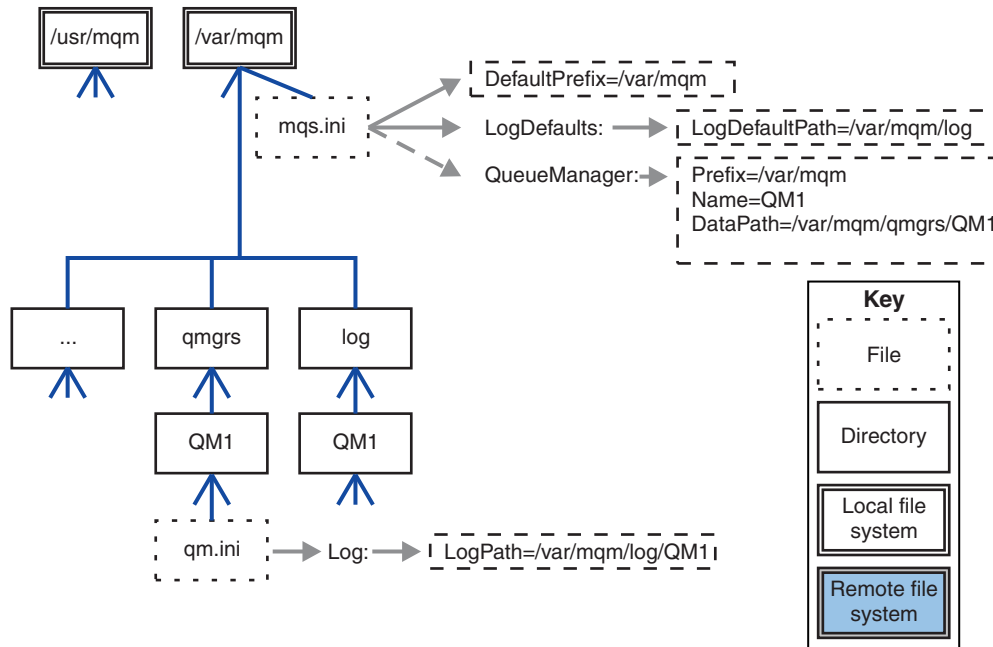


Figure 50. Example default IBM MQ directory structure for UNIX and Linux systems

Typical directory structure for releases earlier than v7.0.1

The `DataPath` attribute did not exist before IBM MQ v7.0.1; the attribute is not present in the `mqs.ini` file. The location of the `qmgrs` directory was configured using the `Prefix` attribute. The location of individual directories could be configured by using symbolic links to point to different file system locations.

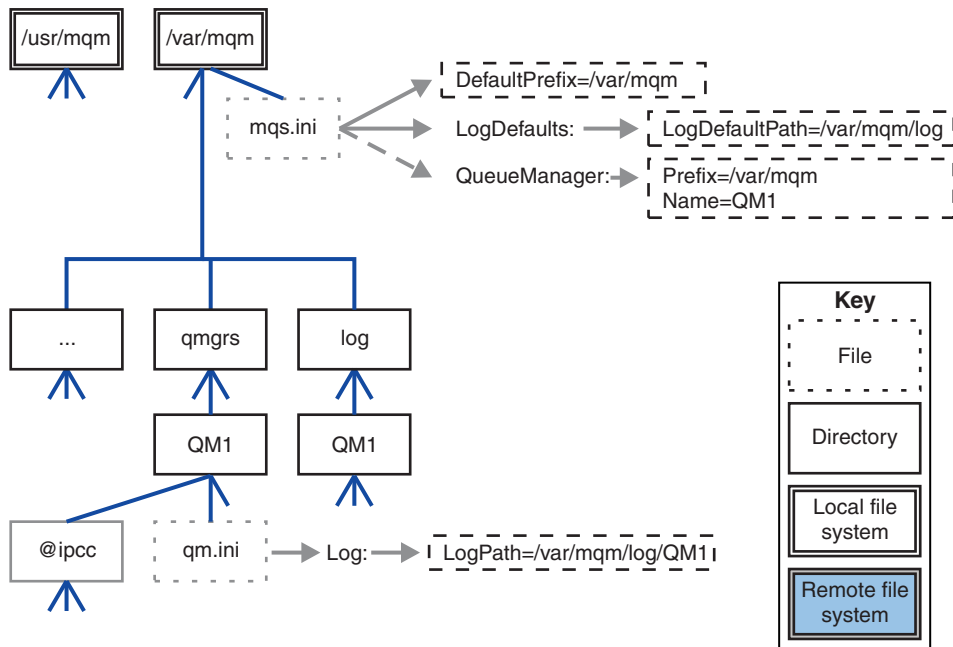


Figure 51. Typical directory structure for releases earlier than v7.0.1

Share default qmgrs and log directories (Release v7.0.1 onwards)

An alternative to “Share everything (Release v7.0.1 onwards)” on page 121, is to share the qmgrs and log directories separately (Figure 52). In this configuration, there is no need to set AMQ_MQS_INI_LOCATION as the default mqs.ini is stored in the local /var/mqm file system. The files and directories, such as mqclient.ini and mqserver.ini are also not shared.

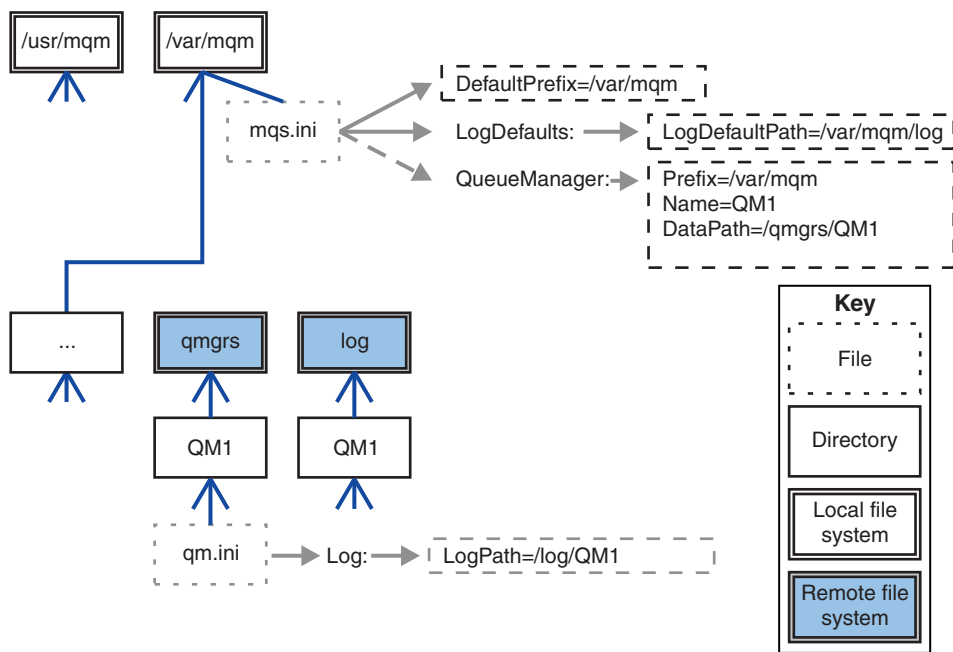


Figure 52. Share qmgrs and log directories

Share named qmgrs and log directories (Release v7.0.1 onwards)

The configuration in Figure 53 places the log and qmgrs in a common named remote shared file system called /ha. The same physical configuration can be created in two different ways.

1. Set LogDefaultPath=/ha and then run the command, **crtmqm - md /ha/qmgrs QM1**. The result is exactly as illustrated in Figure 53.
2. Leave the default paths unchanged and then run the command, **crtmqm - ld /ha/log - md /ha/qmgrs QM1**.

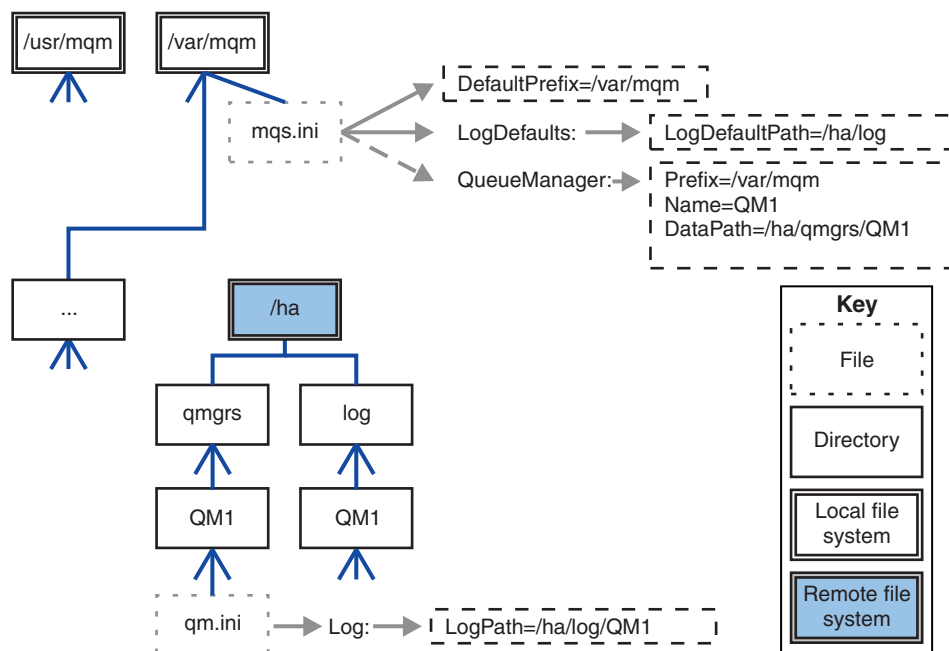


Figure 53. Share named qmgrs and log directories

Share everything (Release v7.0.1 onwards)

Figure 54 on page 122 is a simple configuration for system with fast networked file storage.

Mount /var/mqm as a remote shared file system. By default, when you start QM1, it looks for /var/mqm, finds it on the shared file system, and reads the mqs.ini file in /var/mqm. Rather than use the single /var/mqm/mqs.ini file for queue managers on all your servers, you can set the AMQ_MQS_INI_LOCATION environment variable on each server to point to different mqs.ini files.

Note: The contents of the generic error file in /var/mqm/errors/ are shared between queue managers on different servers.




Figure 54. Share everything

Note that you cannot use this for multi-instance queue managers. The reason is that it is necessary for each host in a multi-instance queue manager to have its own local copy of /var/mqm to keep track of local data, such as semaphores and shared memory. These entities cannot be shared across hosts.

Directory structure on Windows systems

How to find queue manager configuration information and directories on Windows.

The default directory for IBM MQ for Windows installation is:

32 bit C:\Program Files (x86)\WebSphere MQ

64 bit program directories

C:\Program Files\IBM\WebSphere MQ

Data directory

C:\ProgramData\IBM\MQ

Important: For Windows installations, the directories are as stated, unless there is a previous installation of IBM MQ that still contains registry entries or queue managers, or both. In this situation, the new installation uses the old data directory location. For more information, see Program and data directory locations.

A Windows 32-bit client installs on a 32-bit machine, by default, only into C:\Program Files\IBM\WebSphere MQ.

A Windows 64-bit client installs on a 64-bit machine, by default, only into C:\Program Files\IBM\WebSphere MQ.

If you want to know which installation directory and which data directory is being used, run the dspmqver command.

The installation directory is listed in the **InstPath** field and the data directory is listed in the **DataPath** field.

Running the **dspmqver** command displays, for example, the following information:

```
>dspmqver
Name:      WebSphere MQ
Version:   8.0.0.0
Level:     p000-L140303.5
BuildType: IKAP - (Production)
Platform:  IBM MQ for Windows (x64 platform)
Mode:      64-bit
O/S:       Windows 7 Professional x64 Edition, Build 7601: SP1
InstName:  Installation1
InstDesc:
Primary:   Yes
InstPath: C:\Program Files\IBM\WebSphere MQ
DataPath: C:\ProgramData\IBM\MQ
MaxCmdLevel: 800
LicenseType: Production
```

Multi-instance queue managers

To configure a multi-instance queue manager, the log and data directories must be placed on networked storage, preferably on a different server to any of the servers that are running instances of the queue manager.

Two parameters are provided on the **crtmqm** command, **-md** and **-ld**, to make it easier specify the location of the queue manager data and log directories. The effect of specifying the **-md** parameter is fourfold:

1. The `mqs.ini` stanza `QueueManager\QmgrName` contains a new variable, *DataPath*, which points to the queue manager data directory. Unlike the *Prefix* variable, the path includes the name of the queue manager directory.
2. The queue manager configuration information stored in the `mqs.ini` file is reduced to *Name*, *Prefix*, *Directory* and *DataPath*.

Directory content:

Lists the location and content of IBM MQ directories.

An IBM MQ configuration has three main sets of files and directories:

1. Executable, and other read-only files that are only updated when maintenance is applied. For example:
 - The readme file
 - The IBM MQ Explorer plug-in and help files
 - License files

These files are described in Table 14 on page 124.

2. Potentially modifiable files and directories that are not specific to a particular queue manager. These files and directories are described in Table 15 on page 124.
3. Files and directories that are specific to each queue manager on a server. These files and directories are described in Table 16 on page 125.

Resource directories and files

The resource directories and files contain all the executable code and resources to run a queue manager. The variable, *FilePath*, in the installation specific IBM MQ configuration registry key, contains the path to the resource directories.

Table 14. Directories and files in the *FilePath* directory

File path	Contents
<i>FilePath</i> \bin	Commands and DLLs
<i>FilePath</i> \bin64	Commands and DLLs (64 bit)
<i>FilePath</i> \conv	Data conversion tables
<i>FilePath</i> \doc	Wizard help files
<i>FilePath</i> \MQExplorer	Explorer and Explorer help Eclipse plug-ins
<i>FilePath</i> \gskit8	Global security kit
<i>FilePath</i> \java	Java resources, including JRE
<i>FilePath</i> \licenses	License information
<i>FilePath</i> \Non_IBM_License	License information
<i>FilePath</i> \properties	Used internally
<i>FilePath</i> \Tivoli	
<i>FilePath</i> \tools	Development resources and samples
<i>FilePath</i> \Uninst	Used internally
<i>FilePath</i> \README.TXT	Readme file

Directories not specific to a queue manager

Some directories contain files, such as trace files and error logs, that are not specific to a specific queue manager. The *DefaultPrefix* variable contains the path to these directories. *DefaultPrefix* is part of the *AllQueueManagers* stanza.

Table 15. Directories and files in *DefaultPrefix* directory

File path	Contents
<i>DefaultPrefix</i> \Config	Used internally
<i>DefaultPrefix</i> \conv	ccsid.tbl data conversion control file, described in Data conversion
<i>DefaultPrefix</i> \errors	Non queue manager error logs, AMQERR <i>nn</i> .LOG
<i>DefaultPrefix</i> \exits	Channel exit programs
<i>DefaultPrefix</i> \exits64	Channel exit programs (64 bit)
<i>DefaultPrefix</i> \ipc	Not used
<i>DefaultPrefix</i> \Qmgrs	Described in Table 16 on page 125
<i>DefaultPrefix</i> \trace	Trace files
<i>DefaultPrefix</i> \amqmjpse.txt	Used internally

Queue manager directories

When you create a queue manager, a new set of directories, specific to the queue manager, is created.

If you create a queue manager with the **-md** *filepath* parameter, the path is stored in the *DataPath* variable in the queue manager stanza of the *mqs.ini* file. If you create a queue manager without setting the **-md** *filepath* parameter, the queue manager directories are created in the path stored in *DefaultPrefix*, and the path is copied into the *Prefix* variable in the queue manager stanza of the *mqs.ini* file.

Table 16. Directories and files in DataPath and Prefix/Qmgrs/QmgrName directories

File path	Contents	
DataPath\@ipcc	Default location for AMQCLCHL.TAB, the client connection table.	
DataPath\authinfo	Used internally.	
DataPath\channel		
DataPath\cIntconn		
DataPath\errors		Error logs, AMQERR nn.LOG
DataPath\listener	Used internally.	
DataPath\namelist		
DataPath\plugcomp		
DataPath\procdef		
DataPath\qmanager		
DataPath\queues		
DataPath\services		
DataPath\ssl		
DataPath\startprm		
DataPath\topic		
DataPath\active		
DataPath\active.dat		
DataPath\amqalchk.fil		
DataPath\master		
DataPath\master.dat		
DataPath\qm.ini		Queue manager configuration
DataPath\qmstatus.ini		Queue manager status
Prefix\Qmgrs\QmgrName	Used internally	
Prefix\Qmgrs\@SYSTEM	Not used	
Prefix\Qmgrs\@SYSTEM\errors		

Directory structure on IBM i

A description of the IFS is given, and the IBM MQ IFS directory structure is described for server, client, and Java.

The integrated file system (IFS) is a part of IBM i that supports stream input/output and storage management similar to personal computer, UNIX and Linux operating systems, while providing an integrating structure over all information stored in the server.

On IBM i directory names begin with the character & (ampersand) instead of the character @ (at). For example, @system on IBM i is &system.

IFS root file system for IBM MQ server

When you install IBM MQ Server for IBM i, the following directories are created in the IFS root file system.

ProdData:

Overview

```

QIBM  '-- ProdData
        '-- mqm
            '-- doc
            '-- inc
            '-- lib
            '-- samp
            '-- licenses
        '-- LicenseDoc
            '-- 5724H72_V8R0M0

```

/QIBM/ProdData/mqm

Subdirectories below this contain all the product data, for example, C++ classes, trace format files, and license files. Data in this directory is deleted and replaced each time the product is installed.

/QIBM/ProdData/mqm/doc

A Command Reference for the CL commands is provided in HTML format and installed here.

/QIBM/ProdData/mqm/inc

The header files for compiling your C or C++ programs.

/QIBM/ProdData/mqm/lib

Auxiliary files used by MQ.

/QIBM/ProdData/mqm/samp

Further samples.

/QIBM/ProdData/mqm/licenses

License files. The two files for each language are named like LA_ *xx* and LI_ *xx* where *xx* is the 2 character language identifier for each language supplied.

Also the following directory stores license agreements files:

/QIBM/ProdData/LicenseDoc/5724H72_V8R0M0

License files. The files are named like 5724H72_V8R0M0_ *xx* where *xx* is the 2 or 5 character language identifier for each language supplied.

UserData:

Overview

```

QIBM  '-- UserData
        '-- mqm
            '-- errors
            '-- trace
            '-- qmgrs
                '-- &system
                '-- qmgrname1
                '-- qmgrname2
                '-- and so on

```

/QIBM/UserData/mqm

Subdirectories below this contain all user data relating to queue managers.

When you install the product, an mqs.ini file is created in directory /QIBM/UserData/mqm/ (unless it is already there from a previous installation).

When you create a queue manager, a qm.ini file is created in the directory /QIBM/UserData/mqm/qmgrs/ *QMGRNAME* / (where *QMGRNAME* is the name of the queue manager).

Data in the directories is retained when the product is deleted.

IFS root file system for IBM MQ MQI client

When you install IBM MQ MQI client for IBM i, the following directories created in the IFS root file system:

ProdData:

Overview

```
QIBM  '-- ProdData
      '-- mqm
      '-- lib
```

/QIBM/ProdData/mqm

Subdirectories below this directory contain all the product data. Data in this directory is deleted and replaced each time the product is replaced.

UserData:

Overview

```
QIBM  '-- UserData
      '-- mqm
      '-- errors
      '-- trace
```

/QIBM/UserData/mqm

Subdirectories below this directory contain all user data.

IFS root file system for IBM MQ Java

When you install IBM MQ Java on IBM i, the following directories are created in the IFS root file system:

ProdData:

Overview

```
QIBM  '-- ProdData
      '-- mqm
      '-- java
      '-- samples
      '-- bin
      '-- lib
```

/QIBM/ProdData/mqm/java

Subdirectories below this contain all the product data, including Java classes. Data in this directory is deleted and replaced each time the product is replaced.

/QIBM/ProdData/mqm/java/samples

Subdirectories below this contain all the sample Java classes and data.

Libraries created by server and client installations

Installation of the IBM MQ server or client creates the following libraries:

- QMQM

The product library.

- QMQMSAMP

The samples library (if you choose to install the samples).

- QMxxxx

Server only.

Each time that you create a queue manager, IBM MQ automatically creates an associated library, with a name like QMxxxx where xxxx is derived from the queue manager name. This library contains objects specific to the queue manager, including journals and associated receivers. By default the name of this library is derived from the name of the queue manager prefixed with the characters QM. For example, for a queue manager called TEST, the library would be called QMTEST.

Note: When you create a queue manager, you can specify the name of its library if you want to. For example:

```
CRTMQM MQMNAME(TEST) MQMLIB(TESTLIB)
```

You can use the WRKLIB command to list all the libraries that IBM MQ for IBM i has created. Against the queue manager libraries, you will see the text QMGR: QMGRNAME. The format of the command is:

```
WRKLIB LIB(QM*)
```

These queue manager-associated libraries are retained when the product is deleted.

IBM MQ and UNIX System V IPC resources

A queue manager uses some IPC resources. Use **ipcs -a** to find out what resources are being used.

This information applies to IBM MQ running on UNIX and Linux systems only.

IBM MQ uses System V interprocess communication (IPC) resources (*semaphores* and *shared memory segments*) to store and pass data between system components. These resources are used by queue manager processes and applications that connect to the queue manager. IBM MQ MQI clients do not use IPC resources, except for IBM MQ trace control. Use the UNIX command **ipcs -a** to get full information on the number and size of the IPC resources currently in use on the machine.

Shared memory on AIX

If certain application types fail to connect because of an AIX memory limitation, in most cases this can be resolved by setting the environment variable EXTSHM=ON.

Some 32-bit processes on AIX might encounter an operating system limitation that affects their ability to connect to IBM MQ queue managers. Every standard connection to IBM MQ uses shared memory, but unlike other UNIX and Linux platforms, AIX allows 32-bit processes to attach only 11 shared memory sets.

Most 32-bit processes will not encounter this limit, but applications with high memory requirements might fail to connect to IBM MQ with reason code 2102: MQRC_RESOURCE_PROBLEM. The following application types might see this error:

- Programs running in 32-bit Java virtual machines
- Programs using the large or very large memory models
- Programs connecting to many queue managers or databases
- Programs that attach to shared memory sets on their own

AIX offers an extended shared memory feature for 32-bit processes that allows them to attach more shared memory. To run an application with this feature, export the environment variable EXTSHM=ON before starting your queue managers and your program. The EXTSHM=ON feature prevents this error in most cases, but it is incompatible with programs that use the SHM_SIZE option of the shmctl function.

IBM MQ MQI client applications and all 64-bit processes are unaffected by this limitation. They can connect to IBM MQ queue managers regardless of whether EXTSHM has been set.

IBM MQ and UNIX Process Priority

Good practices when setting process priority *nice* values.

This information applies to IBM MQ running on UNIX and Linux systems only.

If you run a process in the background, that process can be given a higher *nice* value (and hence lower priority) by the invoking shell. This might have general IBM MQ performance implications. In highly-stressed situations, if there are many ready-to-run threads at a higher priority and some at a lower priority, operating system scheduling characteristics can deprive the lower priority threads of processor time.

It is good practice that independently started processes associated with queue managers, such as `runmqtsr`, have the same *nice* values as the queue manager they are associated with. Ensure the shell does not assign a higher *nice* value to these background processes. For example, in ksh, use the setting `"set +o bgnice"` to stop ksh from raising the *nice* value of background processes. You can verify the *nice* values of running processes by examining the *NI* column of a `"ps -efl"` listing.

Also, start IBM MQ application processes with the same *nice* value as the queue manager. If they run with different *nice* values, an application thread might block a queue manager thread, or vice versa, causing performance to degrade.

Planning your IBM MQ client environment on HP Integrity NonStop Server

When you are planning your IBM MQ environment, you must consider the HP Integrity NonStop Server environment, and HP NonStop TMF. Use the information to plan the environment where IBM MQ client for HP Integrity NonStop Server runs.

Before you plan your IBM MQ client for HP Integrity NonStop Server architecture, familiarize yourself with the basic IBM MQ client for HP Integrity NonStop Server concepts, see the topics in IBM MQ client for HP Integrity NonStop Server technical overview.

Preparing the HP Integrity NonStop Server environment

Before installation, the environment must be prepared depending on whether the installation is to be verified immediately or not.

For the installation, you require the following items:

- A user ID that meets the requirements. For details about user ID requirements, see [Setting up the user and group on HP Integrity NonStop Server](#).
- Verified locations in the OSS and Guardian file systems that can be for the installation files.
- An operational OSS shell and OSS file system. You can verify the file system by doing the following tasks:
 - Log on to the OSS environment (shell). Ensure that you have write access to the OSS installation root directory you intend to use.
 - Log on to the TAACL environment using the user ID in the MQM group. Verify that the volume you intend to use meets the requirements and is accessible to you, and that the subvolume does not exist.

You can login to both OSS or TAACL using either an alias, if you have one, or your full principal.

If you intend to proceed immediately to verify that the installation is usable, you might also need the following optional items:

- An operational and accessible Local Sockets subsystem in the OSS environment.
- An operational TCP/IP subsystem.

If you intend to use TMF coordinated global units of work, you will need the following items:

- An operational TMF subsystem.
- An operational Pathway (TS/MP) subsystem.

Work with your systems administrator if you are in any doubt about the status of these critical subsystems.

IBM MQ and HP NonStop TMF

IBM MQ client for HP Integrity NonStop Server can participate in HP NonStop Transaction Management Facility (HP NonStop TMF) coordinated units of work. Coordinating transactions with HP NonStop TMF is only supported where the queue manager is at IBM WebSphere MQ Version 7.1 or later.

The IBM MQ provided TMF/Gateway converts transactions from TMF coordination into eXtended Architecture (XA) transaction coordination to communicate with the remote queue manager. The IBM MQ provided TMF/Gateway is the bridge between TMF and queue manager transactions, using the services provided by HP NonStop TMF, and has been designed to run in a Pathway environment.

HP NonStop TMF software provides transaction protection and database consistency in demanding environments. For more information about HP NonStop TMF, see [HP NonStop TMF Introduction](#).

For information about how to configure the IBM MQ provided TMF/Gateway, see [Configuring HP Integrity NonStop Server](#).

Using HP NonStop TMF

The HP NonStop Transaction Management Facility (TMF) is the native transaction manager on HP Integrity NonStop Server and is integrated with the file system and the relational database managers, SQL/MP, and SQL/MX.

IBM MQ client for HP Integrity NonStop Server can use TMF to coordinate global units of work.

To coordinate global units of work, TMF acts as the transaction manager, and an application must use the API provided by TMF to start, commit, and back out global units of work. An application starts a global unit of work by calling `BEGINTRANSACTION`, and then updates IBM MQ resources within the global unit of work by issuing `MQPUT`, `MQPUT1`, and `MQGET` calls within syncpoint control. The application can then commit the global unit of work by calling `ENDTRANSACTION`, or back it out by calling `ABORTTRANSACTION`.

An application that is using TMF transactions can only actively work on one transaction at any one time, however using `RESUMETRANSACTION` allows an application to switch from one active transaction to another, or to being associated with no TMF transaction, without completing or aborting the previously active transaction. Any calls to `MQPUT`, `MQPUT1`, or `MQGET` are made under the currently active TMF transaction, if present, or a local unit of work, if not present. Therefore, care must be taken within the application to ensure that these calls are being made within the correct unit of work.

Within a global unit of work, as well as updating IBM MQ resources, an application can update Enscribe files, SQL/MP databases, or SQL/MX databases.

Using global units of work

A global unit of work is implemented as a TMF transaction. An application starts a global unit of work by calling `BEGINTRANSACTION`, and either commits the unit of work by calling `ENDTRANSACTION` or backs out the unit of work by calling `ABORTTRANSACTION`. An application can use other TMF API calls as well.

An application can inherit a TMF transaction from another application. For example, an application (the first application) can perform work within the transaction before replying and passing the transaction back to a second application for further processing. Both the first and the second applications can therefore participate in the same global unit of work that involves updates to IBM MQ queues and updates to files and databases. The ability to pass a TMF transaction between applications means that several IBM MQ applications can perform messaging operations within the same global unit of work.

An application can manage and control multiple active TMF transactions at the same time. The transactions can be started by the application itself, or inherited from other applications, or both. This means that an application can participate in multiple global units of work at the same time.

The maximum number of concurrent active TMF transactions per process is 1000, which is an architectural limit. If an application is managing multiple TMF transactions, only one transaction can be current at any point in time. Alternatively, none of the transactions can be current. The application can use TMF API calls such as `RESUMETRANSACTION`, `ACTIVATERECEIVETRANSID`, and `TMF_SET_TX_ID` to move the state of being current from one transaction to another, or to designate that no transaction is current. The application uses this level of control to determine whether a messaging operation is performed within a local unit of work, a global unit of work, or outside of syncpoint control:

- If an application calls `MQPUT`, `MQPUT1`, or `MQGET` within syncpoint control when no TMF transaction is current, IBM MQ processes the call within a local unit of work.
- If an application calls `MQPUT`, `MQPUT1`, or `MQGET` within syncpoint control when the application has a current TMF transaction, IBM MQ processes the call within the global unit of work that is implemented by the current TMF transaction.
- If an application calls `MQPUT`, `MQPUT1`, or `MQGET` outside of syncpoint control, IBM MQ processes the call outside of syncpoint control, irrespective of whether the application has a current TMF transaction at the time of the call.

IBM MQ never changes the state of an application's TMF transaction during an MQI call, except when a software or hardware failure occurs during processing and IBM MQ or the operating system determines that the transaction must be backed out to preserve data integrity. Every MQI call restores the transaction state of the application just before returning control to the application.

Avoiding long running transactions

Avoid designing applications in which TMF transactions remain active for more than a few tens of seconds. Long running transactions can cause the circular audit trail of TMF to fill up. Because TMF is a critical system-wide resource, TMF protects itself by backing out application transactions that are active for too long.

Suppose that the processing within an application is driven by getting messages from a queue, and that the application gets a message from the queue and processes the message within a unit of work. Typically, an application calls `MQGET` with the wait option and within syncpoint control to get a message from the queue.

If the application is using a global unit of work instead, the specified wait interval on the `MQGET` call must be short to avoid a long running transaction. This means that the application might need to issue the `MQGET` call more than once before it retrieves a message.

Planning your IBM MQ environment on z/OS

When planning your IBM MQ environment, you must consider the resource requirements for data sets, page sets, Db2, Coupling Facilities, and the need for logging, and backup facilities. Use this topic to plan the environment where IBM MQ runs.

Before you plan your IBM MQ architecture, familiarize yourself with the basic IBM MQ for z/OS concepts, see the topics in IBM MQ for z/OS concepts.

Related concepts:

“Planning” on page 1

When planning your IBM MQ environment, consider the support that IBM MQ provides for single and multiple queue manager architectures, and for point-to-point and publish/subscribe messaging styles. Also plan your resource requirements, and your use of logging and backup facilities.

Related information:

IBM MQ Technical overview

Configuring z/OS

Administering IBM MQ for z/OS

Planning your storage and performance requirements on z/OS

z/OS

You must set realistic and achievable storage, and performance goals for your IBM MQ system. Use this topic help you understand the factors which affect storage, and performance.

This topic contains information about the storage and performance requirements for IBM MQ for z/OS. It contains the following sections:

- z/OS performance options for IBM MQ
- Determining z/OS workload management importance and velocity goals
- “Library storage” on page 133
- “System LX usage” on page 134
- “Address space storage” on page 134
- “Data storage” on page 137

See, “Where to find more information about storage and performance requirements” on page 138 for more information.

z/OS performance options for IBM MQ

With workload management, you define performance goals and assign a business importance to each goal. You define the goals for work in business terms, and the system decides how much resource, such as processor and storage, should be given to the work to meet its goal. Workload management controls the dispatching priority based on the goals you supply. Workload management raises or lowers the priority as needed to meet the specified goal. Thus, you need not fine-tune the exact priorities of every piece of work in the system and can focus instead on business objectives.

The three kinds of goals are:

Response time

How quickly you want the work to be processed

Execution velocity

How fast the work should be run when ready, without being delayed for processor, storage, I/O access, and queue delay

Discretionary

A category for low priority work for which there are no performance goals

Response time goals are appropriate for end-user applications. For example, CICS® users might set workload goals as response time goals. For IBM MQ address spaces, velocity goals are more appropriate. A small amount of the work done in the queue manager is counted toward this velocity goal but this work is critical for performance. Most of the work done by the queue manager counts toward the performance goal of the end-user application. Most of the work done by the channel initiator address space counts toward its own velocity goal. The receiving and sending of IBM MQ messages, which the channel initiator accomplishes, is typically important for the performance of business applications using them.

Determining z/OS workload management importance and velocity goals

For full information about workload management and defining goals through the service definition, see *z/OS MVS Planning: Workload Management*.

This section suggests how to set the z/OS workload management importance and velocity goals relative to other important work in your system.

Use the following service classes:

The default SYSSTC service class

- VTAM and TCP/IP address spaces
- IRLM address space (IRLMPROC)

Note: The VTAM, TCP/IP, and IRLM address spaces must have a higher dispatching priority than all the DBMS address spaces, their attached address spaces, and their subordinate address spaces. Do not allow workload management to reduce the priority of VTAM, TCP/IP, or IRLM to (or below) that of the other DBMS address spaces

A high velocity goal and importance of 1 for a service class with a name that you define, such as PRODREGN, for the following:

- IBM MQ queue manager and channel initiator address spaces
- Db2 (all address spaces, except for the Db2-established stored procedures address space)
- CICS (all region types)
- IMS™ (all region types except BMPs)

A high velocity goal is good for ensuring that startups and restarts are performed as quickly as possible for all these address spaces.

The velocity goals for CICS and IMS regions are only important during startup or restart. After transactions begin running, workload management ignores the CICS or IMS velocity goals and assigns priorities based on the response time goals of the transactions that are running in the regions. These transaction goals should reflect the relative priority of the business applications they implement. They might typically have an importance value of 2. Any batch applications using MQ should similarly have velocity goals and importance reflecting the relative priority of the business applications they implement. Typically the importance and velocity goals will be less than those for PRODREGN.

Library storage

You must allocate storage for the product libraries. The exact figures depend on your configuration, but an estimate of the space required by the distribution libraries is 80 MB. The target libraries require about 72 MB. Additionally, you require space for the SMP/E libraries.

The target libraries used by IBM MQ for z/OS use PDS or PDSE formats. Ensure that any PDSE target libraries are not shared outside a sysplex. For more information about the required libraries and their sizes and the required format, see the Program Directory for WebSphere MQ for z/OS.

System LX usage

Each defined IBM MQ subsystem reserves one system linkage index (LX) at IPL time, and a number of non-system linkage indexes when the queue manager is started. The system linkage index is reused when the queue manager is stopped and restarted. Similarly, distributed queuing reserves one non-system linkage index. In the unlikely event of your z/OS system having inadequate system LXs defined, you might need to take these reserved system LXs into account.

Address space storage

Use this topic for basic guidance on address space requirements for the IBM MQ components.

Storage requirements can be divided into the following categories:

- Common storage
- Queue manager private region storage usage
- Channel initiator storage usage

For more details see, Suggested regions sizes.

With 31-bit address space, a virtual "line" marks the 16-megabyte address, and 31-bit addressable storage is often known as 'above the (16MB) line'. With 64-bit address space there is a second virtual line called 'the bar' that marks the 2-gigabyte address. The bar separates storage below the 2-gigabyte address, called 'below the bar', from storage above the 2-gigabyte address, called 'above the bar'. Storage below the bar uses 31-bit addressability, storage above the bar uses 64-bit addressability.

You can specify the limit of 31-bit storage by using the REGION parameter on the JCL, and the limit of above the bar storage by using the MEMLIMIT parameter. These specified values can be overridden by MVS exits.

Common storage

Each IBM MQ for z/OS subsystem has the following approximate storage requirements:

- CSA 4 KB
- ECSA 800 KB, plus the size of the trace table that is specified in the TRACTBL parameter of the CSQ6SYSP system parameter macro. For more information, see Using CSQ6SYSP.

In addition, each concurrent IBM MQ logical connection requires about 5 KB of ECSA. When a task ends, other IBM MQ tasks can reuse this storage. IBM MQ does not release the storage until the queue manager is shut down, so you can calculate the maximum amount of ECSA required by multiplying the maximum number of concurrent logical connections by 5 KB. Concurrent logical connections are the number of:

- Tasks (TCBs) in Batch, TSO, z/OS UNIX and Linux System Services, IMS, and Db2 SPAS regions that are connected to IBM MQ, but not disconnected.
- CICS transactions that have issued an IBM MQ request, but have not terminated
- JMS Connections, Sessions, TopicSessions or QueueSessions that have been created (for bindings connection), but not yet destroyed or garbage collected.
- Active IBM MQ channels.

You can set a limit to the common storage, used by logical connections to the queue manager, with the ACELIM configuration parameter. The ACELIM control is primarily of interest to sites where Db2 stored procedures cause operations on IBM MQ queues.

When driven from a stored procedure, each IBM MQ operation can result in a new logical connection to the queue manager. Large Db2 units of work, for example due to table load, can result in an excessive demand for common storage.

ACELIM is intended to limit common storage use and to protect the z/OS system. Using ACELIM causes IBM MQ failures when the limit is exceeded. See the ACELIM section in Using CSQ6SYSP for more information.

Use SupportPac MP1B to format the SMF 115 subtype 3 records produced by STATISTICS CLASS(2) trace.

The amount of storage currently in the subpool controlled by the ACELIM value is indicated in the output, on the line titled *ACE/PEB*. SupportPac MP1B indicates the number of bytes in use.

Increase the normal value by a sufficient margin to provide space for growth and workload spikes. Divide the new value by 1024 to yield a maximum storage size in KB for use in the ACELIM configuration.

The channel initiator typically requires ECSA usage of up to 160 KB.

Queue manager private region storage usage

IBM MQ for z/OS can use storage above the 2 GB bar for some internal control blocks. You can have buffer pools in this storage, which gives you the potential to configure much larger buffer pools if sufficient storage is available. Typically buffer pools are the major internal control blocks that use storage above the 2 GB bar.

Each buffer pool size is determined at queue manager initialization time, and storage is allocated for the buffer pool when a page set that is using that buffer pool is connected. A new parameter LOCATION (ABOVE|BELOW) is used to specify where the buffers are allocated. You can use the ALTER BUFFPOOL command to dynamically change the size of buffer pools.

To use above the bar (64 Bit) storage, you can specify a value for MEMLIMIT parameter (for example MEMLIMIT=3G) on the EXEC PGM=CSQYASCP parameter in the queue manager JCL. Your installation might have a default value set.

You should specify a MEMLIMIT and specify a sensible storage size rather than MEMLIMIT=NOLIMIT to prevent potential problems. If you specify NOLIMIT or a very large value, then an ALTER BUFFPOOL command with a large size, can use up all of the available z/OS virtual storage, which will lead to paging in your system.

Start with a MEMLIMIT=3G and increase this size when you need to increase the size of your buffer pools.

Specify MEMLIMIT= 2 GB plus the size of the buffer pools above the bar, rounded up to the nearest GB. For example, for 2 buffer pools configured with LOCATION ABOVE, buffer pool 1 has 10,000 buffers, buffer pool 2 has 50,000 buffers. Memory usage above the bar equals 60,000 (total number of buffers) * 4096 = 245,760,000 bytes = 199.3 MB, rounded up to 200MB as a multiple of 4MB. All buffer pools regardless of LOCATION will make use of 64 bit storage for control structures. As the number of buffer pools and number of buffers in those pools increase this can become significant. A good rule of thumb is that each buffer requires an additional 200 bytes of 64 bit storage. For a configuration with 10 buffer pools each with 20,000 buffers that would require: 200 * 10 * 20,000 = 40,000,000 equivalent to 40 MB. You can specify 3 GB for the MEMLIMIT size, which will allow scope for growth (40MB + 200MB + 2 GB which rounds up to 3 GB).

For some configurations there can be significant performance benefits to using buffer pools that have their buffers permanently backed by real storage. You can achieve this by specifying the FIXED4KB value for the PAGECLAS attribute of the buffer pool. However, you should only do this if there is sufficient real storage available on the LPAR, otherwise other address spaces might be affected. For information about when you should use the FIXED4KB value for PAGECLAS, see IBM MQ Support Pac MP16: IBM MQ for z/OS - Capacity planning & tuning

To minimize paging, consider real storage in addition to the virtual storage that is used by the queue manager and the channel initiator.

Before you use storage above the bar, you should discuss with your MVS systems programmer to ensure that there is sufficient auxiliary storage for peak time usage, and sufficient real storage requirements to prevent paging.

Note:

The size of memory dump data sets might have to be increased to handle the increased virtual storage.

Making the buffer pools so large that there is MVS paging might adversely affect performance. You might consider using a smaller buffer pool that does not page, with IBM MQ moving the message to and from the page set.

You can monitor the address space storage usage from the CSQY220I message that indicates the amount of private region storage in use above and below the 2 GB bar, and the remaining amount.

Channel initiator storage usage

There are two areas of channel initiator storage usage that you must consider:

- Private region
- Accounting and statistics

Private region storage usage

You should specify REGION=0M for the CHINIT to allow it to use the maximum below the bar storage. The storage available to the channel initiator limits the number of concurrent connections the CHINIT can have.

Every channel uses approximately 170 KB of extended private region in the channel initiator address space. Storage is increased by message size if messages larger than 32 KB are transmitted. This increased storage is freed when:

- A sending or client channel requires less than half the current buffer size for 10 consecutive messages.
- A heartbeat is sent or received.

The storage is freed for reuse within the Language Environment, however, is not seen as free by the z/OS virtual storage manager. This means that the upper limit for the number of channels is dependent on message size and arrival patterns, and on limitations of individual user systems on extended private region size. The upper limit on the number of channels is likely to be approximately 9000 on many systems because the extended region size is unlikely to exceed 1.6 GB. The use of message sizes larger than 32 KB reduces the maximum number of channels in the system. For example, if messages that are 100 MB long are transmitted, and an extended region size of 1.6 GB is assumed, the maximum number of channels is 15.

The channel initiator trace is written to a dataspace. The size of the database storage, is controlled by the **TRAXTBL** parameter. See ALTER QMGR.

Accounting and statistics storage usage

You should allow the channel initiator access to a minimum of 256 MB of virtual storage, and you can do this by specifying MEMLIMIT=256M.

If you do not set the MEMLIMIT parameter in the channel initiator JCL, you can set the amount of virtual storage above the bar using the MEMLIMIT parameter in the SMFPRMxx member of SYS1.PARMLIB, or from the IEFUSI exit.

If you set the MEMLIMIT to restrict the above bar storage below the required level, the channel initiator issues the CSQX124E message and class 4 accounting and statistics trace will not be available.

Suggested region sizes

The following table shows suggested values for region sizes.

Table 17. Suggested definitions for JCL region sizes

Definition setting	System
Queue manager	REGION=0M, MEMLIMIT=3G
Channel initiator	REGION=0M

Managing the MEMLIMIT and REGION size

Other mechanisms, for example the MEMLIMIT parameter in the SMFPRMxx member of SYS1.PARMLIB or the IEFUSI exit might be used at your installation to provide a default amount of virtual storage above the bar for z/OS address spaces. See the z/OS MVS Programming: Extended Addressability Guide for full details about limiting storage above the bar.

Data storage

Use this topic when planning your data storage requirements for log data sets, Db2 storage, coupling facility storage, and page data sets.

Work with your storage administrator to determine where to put the queue manager datasets. For example, your storage administrator may give you specific DASD volumes, or SMS storage classes, data classes, and management classes for the different data set types.

- Log data sets must be on DASD. These logs can have high I/O activity with a small response time and do not need to be backed up.
- Archive logs can be on DASD or tape. After they have been created, they might never be read again except in an abnormal situation, such as recovering a page set from a backup. They should have a long retention date.
- Page sets might have low to medium activity and should be backed up regularly. On a high use system, they should be backed up twice a day.
- BSDS datasets should be backed up daily; they do not have high I/O activity.

All datasets are similar to those used by Db2, and similar maintenance procedures can be used for IBM MQ.

See the following sections for details of how to plan your data storage:

- **Logs and archive storage**

“How long do I need to keep archive logs” on page 162 describes how to determine how much storage your active log and archive data sets require, depending on the volume of messages that your IBM MQ system handles and how often the active logs are offloaded to your archive data sets.

- **Db2 storage**

“Db2 storage” on page 155 describes how to determine how much storage Db2 requires for the IBM MQ data.

- **coupling facility storage**

“Defining coupling facility resources” on page 146 describes how to determine how large to make your coupling facility structures.

- **Page set and message storage**

“Planning your page sets and buffer pools” describes how to determine how much storage your page data sets require, depending on the sizes of the messages that your applications exchange, on the numbers of these messages, and on the rate at which they are created or exchanged.

Where to find more information about storage and performance requirements

Use this topic as a reference to find more information about storage and performance requirements.

You can find more information from the following sources:

Table 18. Where to find more information about storage requirements

Topic	Where to look
System parameters	Using CSQ6SYSP and Customizing your queue managers
Storage required to install IBM MQ	Program Directory for WebSphere MQ for z/OS
IEALIMIT and IEFUSI exits	<i>MVS Installation Exits</i> , available from the zSeries website z/OS Internet Library.
Latest information	IBM MQ SupportPac Web site Business Integration - WebSphere MQ SupportPacs.
Workload management and defining goals through the service definition	<i>z/OS MVS Planning: Workload Management</i>

Planning your page sets and buffer pools

Information to help you with planning the initial number, and sizes of your page data sets, and buffer pools.

This topic contains the following sections:

- “Plan your page sets”
 - Page set usage
 - Number of page sets
 - Size of page sets
- “Calculate the size of your page sets” on page 139
 - Page set zero
 - Page set 01 - 99
 - Calculating the storage requirement for messages
- “Enabling dynamic page set expansion” on page 142
- “Defining your buffer pools” on page 143

Plan your page sets

Page set usage

For short-lived messages, few pages are normally used on the page set and there is little or no I/O to the data sets except at startup, during a checkpoint, or at shutdown.

For long-lived messages, those pages containing messages are normally written out to disk. This operation is performed by the queue manager in order to reduce restart time.

Separate short-lived messages from long-lived messages by placing them on different page sets and in different buffer pools.

Number of page sets

Using several large page sets can make the role of the IBM MQ administrator easier because it means that you need fewer page sets, making the mapping of queues to page sets simpler.

Using multiple, smaller page sets has a number of advantages. For example, they take less time to back up, and I/O can be carried out in parallel during backup and restart. However, consider that this adds a significant performance cost to the role of the IBM MQ administrator, who is required to map each queue to one of a much greater number of page sets.

Define at least five page sets, as follows:

- A page set reserved for object definitions (page set zero)
- A page set for system-related messages
- A page set for performance-critical long-lived messages
- A page set for performance-critical short-lived messages
- A page set for all other messages

“Defining your buffer pools” on page 143 explains the performance advantages of distributing your messages on page sets in this way.

Size of page sets

Define sufficient space in your page sets for the expected peak message capacity. Consider for any unexpected peak capacity, such as when a build-up of messages develops because a queue server program is not running. You can do this by allocating the page set with secondary extents or, alternatively, by enabling dynamic page set expansion. For more information, see “Enabling dynamic page set expansion” on page 142.

When planning page set sizes, consider all messages that might be generated, including non-application message data. For example, trigger messages, event messages and any report messages that your application has requested.

The size of the page set determines the time taken to recover a page set when restoring from a backup, because a large page set takes longer to restore.

Note: Recovery of a page set also depends on the time the queue manager takes to process the log records written since the backup was taken; this time period is determined by the backup frequency. For more information, see “Planning for backup and recovery” on page 169.

Note: Page sets larger than 4 GB require the use of SMS extended addressability.

Calculate the size of your page sets

For queue manager object definitions (for example, queues and processes), it is simple to calculate the storage requirement because these objects are of fixed size and are permanent. For messages however, the calculation is more complex for the following reasons:

- Messages vary in size.
- Messages are transitory.
- Space occupied by messages that have been retrieved is reclaimed periodically by an asynchronous process.

Large page sets of greater than 4 GB that provide extra capacity for messages if the network stops, can be created if required. It is not possible to modify the existing page sets. Instead, new page sets with extended addressability and extended format attributes, must be created. The new page sets must be the same physical size as the old ones, and the old page sets must then be copied to the new ones. If backward migration is required, page set zero must not be changed. If page sets less than 4 GB are adequate, no action is needed.

Page set zero

For page set zero, the storage required is:

(maximum number of local queue definitions x 1010)
(excluding shared queues)
+ (maximum number of model queue definitions x 746)
+ (maximum number of alias queue definitions x 338)
+ (maximum number of remote queue definitions x 434)
+ (maximum number of permanent dynamic queue definitions x 1010)
+ (maximum number of process definitions x 674)
+ (maximum number of namelist definitions x 12320)
+ (maximum number of message channel definitions x 2026)
+ (maximum number of client-connection channel definitions x 5170)
+ (maximum number of server-connection channel definitions x 2026)
+ (maximum number of storage class definitions x 266)
+ (maximum number of authentication information definitions x 1010)
+ (maximum number of administrative topic definitions x 15000)
+ (total length of topic strings defined in administrative topic definitions)

Divide this value by 4096 to determine the number of records to specify in the cluster for the page set data set.

You do not need to allow for objects that are stored in the shared repository, but you must allow for objects that are stored or copied to page set zero (objects with a disposition of GROUP or QMGR).

The total number of objects that you can create is limited by the capacity of page set zero. The number of local queues that you can define is limited to 524 287.

Page sets 01 - 99

For page sets 01 - 99, the storage required for each page set is determined by the number and size of the messages stored on that page set. (Messages on shared queues are not stored on page sets.)

Divide this value by 4096 to determine the number of records to specify in the cluster for the page set data set.

Calculating the storage requirement for messages

This section describes how messages are stored on pages. Understanding this can help you calculate how much page set storage you must define for your messages. To calculate the approximate space required for all messages on a page set you must consider maximum queue depth of all the queues that map to the page set and the average size of messages on those queues.

Note: The sizes of the structures and control information given in this section are liable to change between major releases. For details specific to your release of IBM MQ, refer to SupportPac MP16 - WebSphere MQ for z/OS Capacity planning & tuning and MP1E / MP1F / MP1G - WebSphere MQ for z/OS Vx.x.x Performance report

You must allow for the possibility that message "gets" might be delayed for reasons outside the control of IBM MQ (for example, because of a problem with your communications protocol). In this case, the "put" rate of messages might far exceed the "get" rate. This can lead to a large increase in the number of messages stored in the page sets and a consequent increase in the storage size demanded.

Each page in the page set is 4096 bytes long. Allowing for fixed header information, each page has 4057 bytes of space available for storing messages.

When calculating the space required for each message, the first thing you must consider is whether the message fits on one page (a short message) or whether it needs to be split over two or more pages (a long message). When messages are split in this way, you must allow for additional control information in your space calculations.

For the purposes of space calculation, a message can be represented as the following:



The message header section contains the message descriptor and other control information, the size of which varies depending on the size of the message. The message data section contains all the actual message data, and any other headers (for example, the transmission header or the IMS bridge header).

A minimum of two pages are required for page set control information which, is typically less than 1% of the total space required for messages.

Short messages

A short message is defined as a message that fits on one page.

From IBM WebSphere MQ Version 7.0.1, small messages are stored one on each page.

Long messages

If the size of the message data is greater than 3596 bytes, but not greater than 4 MB, the message is classed as a long message. When presented with a long message, IBM MQ stores the message on a series of pages, and stores control information that points to these pages in the same way that it would store a short message. This is shown in Figure 55:

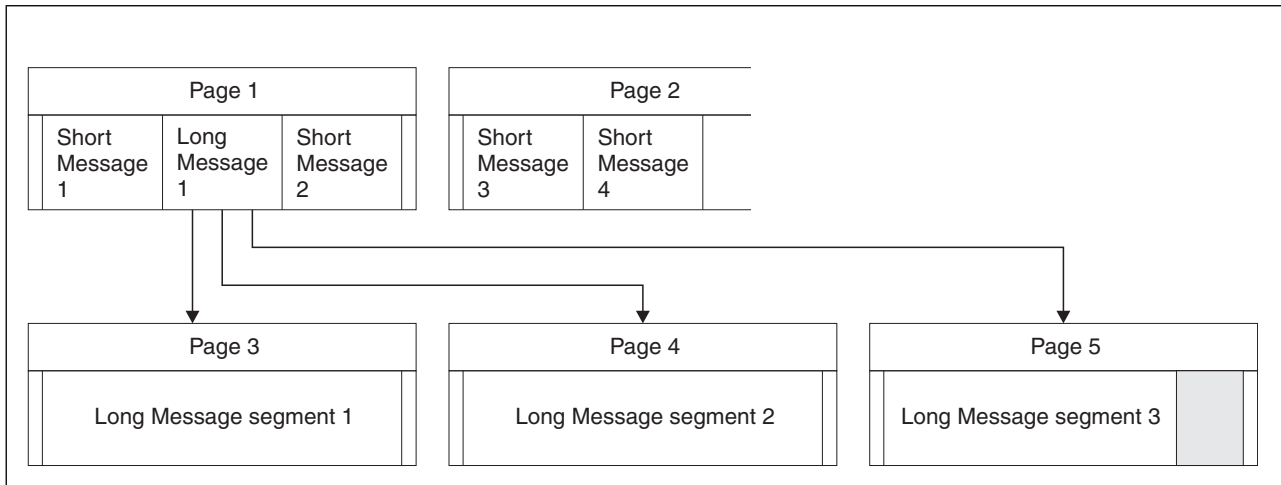


Figure 55. How IBM MQ stores long messages on page sets

Very long messages

Very long messages are messages with a size greater than 4 MB. These are stored so that each 4 MB uses 1037 pages. Any remainder is stored in the same way as a long message, as described above.

Enabling dynamic page set expansion

Page sets can be extended dynamically while the queue manager is running. A page set can have 119 extents, and can be spread over multiple disk volumes.

Each time a page set expands, a new data set extent is used. The queue manager continues to expand a page set when required, until the maximum number of extents has been reached, or until no more storage is available for allocation on eligible volumes.

Once page set expansion fails for one of the reasons above, the queue manager marks the page set for no further expansion attempts. This marking can be reset by altering the page set to EXPAND(SYSTEM).

Page set expansion takes place asynchronously to all other page set activity, when 90% of the existing space in the page set is allocated.

The page set expansion process formats the newly allocated extent and makes it available for use by the queue manager. However, none of the space is available for use, until the entire extent has been formatted. This means that expansion by a large extent is likely to take some time, and putting applications might 'block' if they fill the remaining 10% of the page set before the expansion has completed.

Sample thlqual.SCSQPROC(CSQ4PAGE) shows how to define the secondary extents.

To control the size of new extents, you use one of the following options of the EXPAND keyword of the DEFINE PSID and ALTER PSID commands:

- USER
- SYSTEM
- NONE

USER

Uses the secondary extent size specified when the page set was allocated. If a value was not specified, or if a value of zero was specified, dynamic page set expansion cannot occur.

Page set expansion occurs when the space in the page is 90% used, and is performed asynchronously with other page set activity.

This may lead to expansion by more than a single extent at a time.

Consider the following example: you allocate a page set with a primary extent of 100000 pages and a secondary extent of 5000 pages. A message is put that requires 9999 pages. If the pageset is already using 85,000 pages, writing the message crosses the 90% full boundary (90,000 pages). At this point, a further secondary extent is allocated to the primary extent of 100,000 pages, taking the page set size to 105,000 pages. The remaining 4999 pages of the message continue to be written. When the used page space reaches 94,500 pages, which is 90% of the updated page set size of 105,000 pages, another 5000 page extent is allocated, taking the page set size to 110,000 pages. At the end of the MQPUT, the pageset has expanded twice, and 94,500 pages are used. None of the pages in the second page set expansion have been used, although they were allocated.

At restart, if a previously used page set has been replaced with a data set that is smaller, it is expanded until it reaches the size of the previously used data set. Only one extent is required to reach this size.

SYSTEM

Ignores the secondary extent size that was specified when the page set was defined. Instead, the queue manager sets a value that is approximately 10% of the current page set size. The value is rounded up to the nearest cylinder of DASD.

If a value was not specified, or if a value of zero was specified, dynamic page set expansion can still occur. The queue manager sets a value that is approximately 10% of the current page set size. The new value is rounded up depending on the characteristics of the DASD.

Page set expansion occurs when the space in the page set is approximately 90% used, and is performed asynchronously with other page set activity.

At restart, if a previously used page set has been replaced with a data set that is smaller, it is expanded until it reaches the size of the previously used data set.

NONE

No further page set expansion is to take place.

Related information:

ALTER PSID

DEFINE PSID

DISPLAY USAGE

Defining your buffer pools

Use this topic to help plan the number of buffer pools you should define, and their settings.

This topic is divided into the following sections:

1. "Decide on the number of buffer pools to define"
2. "Decide on the initial characteristics of each buffer pool" on page 144
3. "Monitor the performance of buffer pools under expected load" on page 145
4. "Adjust buffer pool characteristics" on page 145

Decide on the number of buffer pools to define

You should define four buffer pools initially:

Buffer pool 0

Use for object definitions (in page set zero) and performance critical, system related message queues, such as the SYSTEM.CHANNEL.SYNCQ queue and the SYSTEM.CLUSTER.* queues.

Use the remaining three buffer pools for user messages.

Buffer pool 1

Use for important long-lived messages.

Long-lived messages are those that remain in the system for longer than two checkpoints, at which time they are written out to the page set. If you have many long-lived messages, this buffer pool should be relatively small, so that page set I/O is evenly distributed (older messages are written out to DASD each time the buffer pool becomes 85% full).

If the buffer pool is too large, and the buffer pool never gets to 85% full, page set I/O is delayed until checkpoint processing. This might affect response times throughout the system.

If you expect few long-lived messages only, define this buffer pool so that it is sufficiently large to hold all these messages.

Buffer pool 2

Use for performance-critical, short-lived messages.

There is normally a high degree of buffer reuse, using few buffers. However, you should make this buffer pool large to allow for unexpected message accumulation, for example, when a server application fails.

Buffer pool 3

Use for all other (typically, performance noncritical) messages.

Queues such as the dead-letter queue, SYSTEM.COMMAND.* queues and SYSTEM.ADMIN.* queues can also be mapped to buffer pool 3.

Where virtual storage constraints exist, and buffer pools need to be smaller, buffer pool 3 is the first candidate for size reduction.

You might need to define additional buffer pools in the following circumstances:

- If a particular queue is known to require isolation, perhaps because it exhibits different behavior at various times.
 - Such a queue might either require the best performance possible under the varying circumstances, or need to be isolated so that it does not adversely affect the other queues in a buffer pool.
 - Each such queue can be isolated into its own buffer pool and pageset.
- You want to isolate different sets of queues from each other for class-of-service reasons.
 - Each set of queues might then require one, or both, of the two types of buffer pools 1 or 2, as described in Table 20 on page 145, necessitating creation of several buffer pools of a specific type.

In IBM MQ Version 8.0 for z/OS, the maximum number of buffer pools that you can define depends on the OPMODE parameter. If you specify:

- OPMODE(COMPAT,800), or equivalent, a maximum of 16 buffer pools can be defined
- OPMODE(NEWFUNC,800) a maximum of 100 buffer pools can be defined.

Decide on the initial characteristics of each buffer pool

Having decided on the number of buffer pools that you require, you now need to decide the initial characteristics of those buffer pools. The available characteristics are affected by OPMODE, and are summarized in Table 19.

Table 19. Buffer pool characteristics by OPMODE setting

Parameter name/ OPMODE setting	Maximum number of buffers in a single buffer pool	Maximum number of buffer pools	Maximum total number of buffers in queue manager	Buffer pool location	Buffer pool page class
DEFINE or ../ com.ibm.mq.ref.adm.doc/ q085150_.dita parameter	BUFFERS			LOCATION (ABOVE/ BELOW)	PAGECLAS (4KB/FIXED4KB)
OPMODE (COMPAT,800) or equivalent	500 000	16	Limited by available space below the bar. Likely to be less than 262 144	BELOW bar	Pageable (4KB)
OPMODE (NEWFUNC, 800)	999 999 999	100	99 999 999 900 (If MEMLIMIT /IEFUSI exit allows)	ABOVE or BELOW bar	Pageable (4KB) or page-fixed (FIXED4KB)

If you are using the four buffer pools described in “Decide on the number of buffer pools to define” on page 143, then Table 20 on page 145 gives two sets of values for the size of the buffer pools.

The first set is suitable for a test system, the other for a production system or a system that will become a production system eventually. Regardless of the OPMODE setting, the values assume that the buffer pools will be located below the bar, and that the buffers will be pageable.

Table 20. Suggested definitions for buffer pool settings

Definition setting	Test system	Production system
BUFFPOOL 0	1 050 buffers	50 000 buffers
BUFFPOOL 1	1 050 buffers	20 000 buffers
BUFFPOOL 2	1 050 buffers	50 000 buffers
BUFFPOOL 3	1 050 buffers	20 000 buffers

If you need more than the four suggested buffer pools, select the buffer pool (1 or 2) that most accurately describes the expected behavior of the queues in the buffer pool, and size it using the information in Table 20.

It might be that you will need to reduce the size of some of the other buffer pools, or reconsider the number of buffer pools, especially if you have specified OPMODE(COMPAT,800), or equivalent.

Monitor the performance of buffer pools under expected load

You can monitor the usage of buffer pools by analyzing buffer pool performance statistics. In particular, you should ensure that the buffer pools are large enough so that the values of QPSTSOS, QPSTSTLA, and QPSTDMC remain at zero.

For further information, see Buffer manager data records.

Adjust buffer pool characteristics

Use the following points to adjust the buffer pool settings from “Decide on the initial characteristics of each buffer pool” on page 144, if required.

Use the performance statistics from “Monitor the performance of buffer pools under expected load” as guidance.

Note: These points all assume that you have specified OPMODE(NEWFUNC,800), with the exception of point 2.

1. If you are migrating from an earlier version of IBM MQ, only change your existing settings if you have more real storage available.
2. In general, bigger buffer pools are better for performance, and buffer pools can be much bigger if they are above the bar.

However, at all times you should have sufficient real storage available so that the buffer pools are resident in real storage. It is better to have smaller buffer pools that do not result in paging, than big ones that do.

Additionally, there is no point having a buffer pool that is bigger than the total size of the pagesets that use it, although you should take into account pageset expansion if it is likely to occur.

3. Aim for one buffer pool per pageset, as this provides better application isolation.
4. If you have sufficient real storage, such that your buffer pools will never be paged out by the operating system, consider using page-fixed buffers in your buffer pool.
This is particularly important if the buffer pool is likely to undergo much I/O, as it saves the CPU cost associated with page-fixing the buffers before the I/O, and page-unfixing them afterwards.
5. There are several benefits to locating buffer pools above the bar even if they are small enough to fit below the bar. These are:

- 31 bit virtual storage constraint relief - for example more space for common storage.

- If the size of a buffer pool needs to be increased unexpectedly while it is being heavily used, there is less impact and risk to the queue manager, and its workload, by adding more buffers to a buffer pool that is already above the bar, than moving the buffer pool to above the bar and then adding more buffers.
6. Tune buffer pool zero and the buffer pool for short-lived messages (buffer pool 2) so that the 15% free threshold is never exceeded (that is, QPSTCBSL divided by QPSTNBUF is always greater than 15%). If more than 15% of buffers remain free, I/O to the page sets using these buffer pools can be largely avoided during normal operation, although messages older than two checkpoints are written to page sets.
Attention: The optimum value for these parameters is dependent on the characteristics of the individual system. The values given are intended only as a guideline and might not be appropriate for your system.
 7. SYSTEM.* queues which get very deep, for example SYSTEM.CHANNEL.SYNCQ, might benefit from being placed in their own buffer pool, if sufficient storage is available.

IBM MQ SupportPac MP16 - WebSphere MQ for z/OS Capacity planning & tuning provides further information about tuning buffer pools.

Planning your coupling facility and offload storage environment

Use this topic when planning the initial sizes, and formats of your coupling facility (CF) structures, and shared message data set (SMDS) environment or Db2 environment.

This section contains information about the following topics:

- “Defining coupling facility resources”
 - Deciding your offload storage mechanism
 - Planning your structures
 - Planning the size of your structures
 - Mapping shared queues to structures
- “Planning your shared message data set (SMDS) environment” on page 151
- “Planning your Db2 environment” on page 155

Defining coupling facility resources

If you intend to use shared queues, you must define the coupling facility structures that IBM MQ will use in your CFRM policy. To do this you must first update your CFRM policy with information about the structures, and then activate the policy.

Your installation probably has an existing CFRM policy that describes the Coupling Facilities available. The IXCMIAPU z/OS utility is used to modify the contents of the policy based on textual statements you provide. The utility is described in the *MVS Setting up a Sysplex* manual. You must add statements to the policy that define the names of the new structures, the Coupling Facilities that they are defined in, and what size the structures are.

The CFRM policy also determines whether IBM MQ structures are duplexed and how they are reallocated in failure scenarios. Shared queue recovery contains recommendations for configuring CFRM for System Managed Rebuild processing.

Deciding your offload storage environment

The message data for shared queues can be offloaded from the coupling facility and stored in either a Db2 table or in an IBM MQ managed data set called a *shared message data set* (SMDS). Messages which are too large to store in the coupling facility (that is, larger than 63 KB) must always be offloaded, and smaller messages may optionally be offloaded to reduce coupling facility space usage.

For more information, see Specifying offload options for shared messages.

Planning your structures

A queue-sharing group requires a minimum of two structures to be defined. The first structure, known as the administrative structure, is used to coordinate IBM MQ internal activity across the queue-sharing group. No user data is held in this structure. It has a fixed name of *qsg-name* CSQ_ADMIN (where *qsg-name* is the name of your queue-sharing group). Subsequent structures are used to hold the messages on IBM MQ shared queues. Each structure can hold up to 512 shared queues.

Using multiple structures

A queue-sharing group can connect to up to 64 coupling facility structures. One of these structures must be the administration structure, one of these structures might be the SYSAPPL structure. So you can use up to 63 (62 with SYSAPPL) structures for IBM MQ data. You might choose to use multiple structures for any of the following reasons:

- You have some queues that are likely to hold a large number of messages and so require all the resources of an entire coupling facility.
- You have a requirement for a large number of shared queues, so they must be split across multiple structures because each structure can contain only 512 queues.
- RMF reports on the usage characteristic of a structure suggest that you should distribute the queues it contains across a number of Coupling Facilities.
- You want some queue data to held in a physically different coupling facility from other queue data for data isolation reasons.
- Recovery of persistent shared messages is performed using structure level attributes and commands, for example BACKUP CFSTRUCT. To simplify backup and recovery, you could assign queues that hold nonpersistent messages to different structures from those structures that hold persistent messages.

When choosing which Coupling Facilities to allocate the structures in, consider the following points:

- Your data isolation requirements.
- The volatility of the coupling facility (that is, its ability to preserve data through a power outage).
- Failure independence between the accessing systems and the coupling facility, or between Coupling Facilities.
- The level of coupling facility Control Code (CFCC) installed on the coupling facility (IBM MQ requires Level 9 or higher).

Planning the size of your structures

The administrative structure (*qsg-name* CSQ_ADMIN) must be large enough to contain 1000 list entries for each queue manager in the queue-sharing group. When a queue manager starts, the structure is checked to see if it is large enough for the number of queue managers currently *defined* to the queue-sharing group. Queue managers are considered as being defined to the queue-sharing group if they have been added by the CSQ5PQSG utility. You can check which queue managers are defined to the group with the MQSC DISPLAY GROUP command.

Table 21 on page 148 shows the minimum required size for the administrative structure for various numbers of queue managers defined in the queue-sharing group. These sizes were established for a CFCC level 14 coupling facility structure; for higher levels of CFCC, they probably need to be larger.

Table 21. Minimum administrative structure sizes

Number of queue managers defined in queue-sharing group	Required storage
1	6144 KB
2	6912 KB
3	7976 KB
4	8704 KB
5	9728 KB
6	10496 KB
7	11520 KB
8	12288 KB
9	13056 KB
10	14080 KB
11	14848 KB
12	15616 KB
13	16640 KB
14	17408 KB
15	18176 KB
16	19200 KB
17	19968 KB
18	20736 KB
19	21760 KB
20	22528 KB
21	23296 KB
22	24320 KB
23	25088 KB
24	25856 KB
25	27136 KB
26	27904 KB
27	28672 KB
28	29696 KB
29	30464 KB
30	31232 KB
31	32256 KB

When you add a queue manager to an existing queue-sharing group, the storage requirement might have increased beyond the size recommended in Table 21. If so, use the following procedure to estimate the required storage for the CSQ_ADMIN structure: Issue MQSC command /pf DISPLAY CFSTATUS(*), where /cpf is for an existing member of the queue-sharing group, and extract the ENTSMAX information for the CSQ_ADMIN structure. If this number is less than 1000 times the total number of queue managers you want to define in the queue-sharing group (as reported by the DISPLAY GROUP command), increase the structure size.

The size of the structures required to hold IBM MQ messages depends on the likely number and size of the messages to be held on a structure concurrently, together with an estimate of the likely number of concurrent units of work.

The graph in Figure 56 shows how large you should make your CF structures to hold the messages on your shared queues. To calculate the allocation size you need to know

- The average size of messages on your queues
- The total number of messages likely to be stored in the structure

Find the number of messages along the horizontal axis. (Ticks are at multiples of 2, 5, and 8.) Select the curve that corresponds to your message size and determine the required value from the vertical axis. For example, for 200 000 messages of length 1 KB gives a value in the range 256 through 512MB.

Table 22 on page 150 provides the same information in tabular form.




Figure 56. Calculating the size of a coupling facility structure

Use this table to help calculate how large to make your coupling facility structures:

Table 22. Calculating the size of a coupling facility structure

Number of messages	1 KB	2 KB	4 KB	8 KB	16 KB	32 KB	63 KB
100	6	6	7	7	8	10	14
1000	8	9	12	17	27	48	88
10000	25	38	64	115	218	423	821
100000	199	327	584	1097	2124	4177	8156

Your CFRM policy should include the following statements:

INITSIZE is the size in KB that XES allocates to the structure when the first connector connects to it. SIZE is the maximum size that the structure can attain. FULLTHRESHOLD sets the percentage value of the threshold at which XES issues message IXC585E to indicate that the structure is getting full. A best practice is to ensure that INITSIZE and SIZE are within a factor of 2.

For example, with the figures determined previously, you might include the following statements:

```
STRUCTURE NAME(structure-name)
INITSIZE(value from graph in KB, that is, multiplied by 1024)
SIZE(something larger)
FULLTHRESHOLD(85)

STRUCTURE NAME(QSG1APPLICATION1)
INITSIZE(272144) /* 256 MB */
SIZE(524288) /* 512 MB */
FULLTHRESHOLD(85)
```

If the structure use reaches the threshold where warning messages are issued, intervention is required. You might use IBM MQ to inhibit **MQPUT** operations to some of the queues in the structure to prevent applications from writing more messages, start more applications to get messages from the queues, or quiesce some of the applications that are putting messages to the queue.

Alternatively, you can use XES facilities to alter the structure size in place. The following z/OS command:
SETXCF START,ALTER,STRNAME= *structure-name*,SIZE= *newsiz*e

alters the size of the structure to *newsiz*e, where *newsiz*e is a value that is less than the value of SIZE specified on the CFRM policy for the structure, but greater than the current coupling facility size.

You can monitor the use of a coupling facility structure with the MQSC DISPLAY GROUP command.

If no action is taken and a queue structure fills up, an MQRC_STORAGE_MEDIUM_FULL return code is returned to the application. If the administration structure becomes full, the exact symptoms depend on which processes experience the error, but they might include the following problems:

- No responses to commands.
- Queue manager failure as a result of problems during commit processing.

Starting in V7.0.1 certain system queues are provided with CFSTRUCT attributes which specify an application structure CSQSYSAPPL prefixed with the queue-sharing group name. The CSQSYSAPPL structure is an application structure for system queues. For details of creating the coupling facility structures see Task 10: Set up the coupling facility.

With the default definitions, the SYSTEM.QSG.CHANNEL.SYNCQ and SYSTEM.QSG.UR.RESOLUTION.QUEUE use this structure. Table 3 demonstrates an example of how to estimate the message data sizes for the default queues.

Table 23. Table showing CSQSYSAPPL usage against sizing.

<i>qsg-name</i> CSQSYSAPPL usage	sizing
SYSTEM.QSG.CHANNEL.SYNCQ	2 messages of 500 bytes per active instance of a shared channel
SYSTEM.QSG.UR.RESOLUTION.QUEUE	1000 messages of 2 KB

The suggested initial structure definition values are as follows:

```
STRUCTURE NAME(qsgname CSQSYSAPPL)
INITSIZE(20480) /* 20 MB */
SIZE(30720) /* 30 MB */
FULLTHRESHOLD(85)
```

These values can be adjusted depending on your use of shared channels and group units of recovery.

Mapping shared queues to structures

The CFSTRUCT attribute of the queue definition is used to map the queue to a structure.

IBM MQ adds the name of the queue-sharing group to the beginning of the CFSTRUCT attribute. For a structure defined in the CFRM policy with name *qsg-name* SHAREDQ01, the definition of a queue that uses this structure is:

```
DEFINE QLOCAL( myqueue ) QSGDISP(SHARED) CFSTRUCT(SHAREDQ01)
```

Planning your shared message data set (SMDS) environment

If you are using queue-sharing groups with SMDS offloading, IBM MQ needs to connect to a group of shared message data sets. Use this topic to help understand the data set requirements, and configuration required to store IBM MQ message data.

A *shared message data set* (described by the keyword SMDS) is a data set used by a queue manager to store offloaded message data for shared messages stored in a coupling facility structure.

When this form of data offloading is enabled, the **CFSTRUCT** requires an associated group of shared message data sets, one data set for each queue manager in the queue-sharing group. The group of shared message data sets is defined to IBM MQ using the **DSGROUP** parameter on the **CFSTRUCT** definition. Additional parameters can be used to supply further optional information, such as the number of buffers to use and expansion attributes for the data sets.

Each queue manager can write to the data set which it owns, to store shared message data for messages written via that queue manager, and can read all of the data sets in the group

A list describing the status and attributes for each data set associated with the structure is maintained internally as part of the **CFSTRUCT** definition, so each queue manager can check the definition to find out which data sets are currently available.

This data set information can be displayed using the **DISPLAY CFSTATUS TYPE(SMDS)** command to display current status and availability, and the **DISPLAY SMDS** command to display the parameter settings for the data sets associated with a specified **CFSTRUCT**.

Individual shared message data sets are effectively identified by the combination of the owning queue manager name (usually specified using the **SMDS** keyword) and the **CFSTRUCT** structure name.

This section describes the following topics:

- The DSGROUP parameter
- The DSBLOCK parameter

- Shared message data set characteristics
- Shared message data set space management
- Access to shared message data sets
- Creating a shared message data set
- Shared message data set performance and capacity considerations
- Activating a shared message data set

See DEFINE CFSTRUCT for details of these parameters.

For information on managing your shared message data sets, see Managing shared message data sets for further details.

The DSGROUP parameter

The **DSGROUP** parameter on the **CFSTRUCT** definition identifies the group of data sets in which large messages for that structure are to be stored. Additional parameters may be used to specify the logical block size to be used for space allocation purposes and values for the buffer pool size and automatic data set expansion options.

The **DSGROUP** parameter must be set up before offloading to data sets can be enabled.

- If a new **CFSTRUCT** is being defined at **CFLEVEL(5)** and the option **OFFLOAD(SMDS)** is specified or assumed, then the **DSGROUP** parameter must be specified on the same command.
- If an existing **CFSTRUCT** is being altered to increase the **CFLEVEL** to **CFLEVEL(5)** and the option **OFFLOAD(SMDS)** is specified or assumed, then the **DSGROUP** parameter must be specified on the same command if it is not already set.

The DSBLOCK parameter

Space within each data set is allocated to queues as logical blocks of a fixed size (usually 256 KB) specified using the **DSBLOCK** parameter on the **CFSTRUCT** definition, then allocated to individual messages as ranges of pages of 4 KB (corresponding to the physical block size and control interval size) within each logical block. The logical block size also determines the maximum amount of message data that can be read or written in a single I/O operation, which is the same as the buffer size for the SMDS buffer pool.

A larger value of the **DSBLOCK** parameter can improve performance for very large messages by reducing the number of separate I/O operations. However, a smaller value decreases the amount of buffer storage required for each active request. The default value for the **DSBLOCK** parameter is 256 KB, which provides a reasonable balance between these requirements, so specifying this parameter might not normally be necessary.

Shared message data set characteristics

A shared message data set is defined as a VSAM linear data set (LDS). Each offloaded message is stored in one or more blocks in the data set. The stored data is addressed directly by information in the coupling facility entries, like an extended form of virtual storage. There is no separate index or similar control information stored in the data set itself.

The direct addressing scheme means that for messages which fit into one block, only a single I/O operation is needed to read or write the block. When a message spans more than one block, the I/O operations for each block can be fully overlapped to minimize elapsed time, provided that sufficient buffers are available.

The shared message data set also contains a small amount of general control information, consisting of a header in the first page, which includes recovery and restart status information, and a space map

checkpoint area which is used to save the free block space map at queue manager normal termination.

Shared message data set space management

As background information for capacity, performance and operational considerations, it might be useful to understand the concepts of how space in shared message data sets is managed by the queue managers.

Free space in each shared message data set is tracked by its owning queue manager using a space map which indicates the number of pages in use within each logical block. The space map is maintained in main storage while the data set is open and saved in the data set when it is closed normally. (In recovery situations the space map is automatically rebuilt by scanning the messages in the coupling facility structure to find out which data set pages are currently in use).

When a shared message with offloaded message data is being written, the queue manager allocates a range of pages for each message block. If there is a partly used current logical block for the specified queue, the queue manager allocates space starting at the next free page in that block, otherwise it allocates a new logical block. If the whole message does not fit within the current logical block, the queue manager splits the message data at the end of the logical block and allocates a new logical block for the next message block. This is repeated until space has been allocated for the whole message. Any unused space in the last logical block is saved as the new current logical block for the queue. When the data set is closed normally, any unused pages in current logical blocks are returned to the space map before it is saved.

When a shared message with offloaded message data has been read and is ready to be deleted, the queue manager processes the delete request by transferring the coupling facility entry for the message to a clean-up list monitored by the owning queue manager (which may be the same queue manager). When entries arrive on this list, the owning queue manager reads and deletes the entries and returns the freed ranges of pages to the space map. When all used pages in a logical block have been freed the block becomes available for reuse.

Access to shared message data sets

Each shared message data set must be on shared direct access storage which is accessible to all queue managers in the queue-sharing group.

During normal running, each queue manager opens its own shared message data set for read/write access, and opens any active shared message data sets for other queue managers for read-only access, so it can read messages stored by those queue managers. This means that each queue manager userid requires at least UPDATE access to its own shared message data set and READ access to all other shared message data sets for the structure.

If it is necessary to recover shared message data sets using **RECOVER CFSTRUCT**, the recovery process can be executed from any queue manager in the queue-sharing group. A queue manager which may be used to perform recovery processing requires UPDATE access to all data sets that it may need to recover

Creating a shared message data set

Each shared message data set should normally be created before the corresponding **CFSTRUCT** definition is created or altered to enable the use of this form of message offloading, as the **CFSTRUCT** definition changes will normally take effect immediately, and the data set will be required as soon as a queue manager attempts to access a shared queue which has been assigned to that structure. A sample job to allocate and pre-format a shared message data set is provided in `SCSQPROC(CSQ4SMDS)`. The job must be customized and run to allocate a shared message data set for each queue manager which uses a **CFSTRUCT** with **OFFLOAD(SMDS)**.

If the queue manager finds that offload support has been enabled and tries to open its shared message data set but it has not yet been created, the shared message data set will be flagged as unavailable. The queue manager will then be unable to store any large messages until the data set has been created and the queue manager has been notified to try again, for example using the **START SMDSCONN** command.

A shared message data set is created as a VSAM linear data set using an Access Method Services **DEFINE CLUSTER** command. The definition must specify **SHAREOPTIONS(2 3)** to allow one queue manager to open it for write access and any number of queue managers to read it at the same time. The default control interval size of 4 KB must be used. If the data set may need to expand beyond 4 GB, it must be defined using an SMS data class which has the VSAM extended addressability attribute.

Each shared message data set can either be empty or pre-formatted to binary zeros (using **CSQJUFMT** or a similar utility such as the sample job **SCSQPROC(CSQ4SMDS)**), before its initial use. If it is empty or only partly formatted when it is opened, the queue manager automatically formats the remaining space to binary zeros.

Shared message data set performance and capacity considerations

Each shared message data set is used to store offloaded data for shared messages written to the associated **CFSTRUCT** by the owning queue manager, from regions within the same system. The stored data for each message includes a descriptor (currently about 350 bytes), the message headers and the message body. Each offloaded message is stored in one or more pages (physical blocks of size 4 KB) in the data set.

The data set space required for a given number of offloaded messages can therefore be estimated by rounding up the overall message size (including the descriptor) to the next multiple of 4 KB and then multiplying by the number of messages.

As for a page set, when a shared message data set is almost full, it can optionally be expanded automatically. The default behavior for this automatic expansion can be set using the **DSEXPA** parameter on the **CFSTRUCT** definition. This setting can be overridden for each queue manager using the **DSEXPA** parameter on the **ALTER SMDS** command. Automatic expansion is triggered when the data set reaches 90% full and more space is required. If expansion is allowed but an expansion attempt is rejected by VSAM because no secondary space allocation was specified when the data set was defined, expansion is retried using a secondary allocation of 20% of the current size of the data set.

Provided that the shared message data set is defined with the extended addressability attribute, the maximum size is only limited by VSAM considerations to a maximum of 16 TB or 59 volumes. This is significantly larger than the 64 GB maximum size of a local page set.

Activating a shared message data set

When a queue manager has successfully connected to an application coupling facility structure, it checks whether that structure definition specifies offloading using an associated **DSGROUP** parameter. If so, the queue manager allocates and opens its own shared message data set for write access, then it opens for read access any existing shared message data sets owned by other queue managers.

When a shared message data set is opened for the first time (before it has been recorded as active within the queue-sharing group), the first page will not yet contain a valid header. The queue manager fills in header information to identify the queue-sharing group, the structure name and the owning queue manager.

After the header has been completed, the queue manager registers the new shared message data set as active and broadcasts an event to notify any other active queue managers about the new data set.

Every time a queue manager opens a shared message data set it validates the header information to ensure that the correct data set is still being used and that it has not been damaged.

Planning your Db2 environment

If you are using queue-sharing groups, IBM MQ needs to attach to a Db2 subsystem that is a member of a data sharing group. Use this topic to help understand the Db2 requirements used to hold IBM MQ data.

IBM MQ needs to know the name of the data sharing group that it is to connect to, and the name of a Db2 subsystem (or Db2 group) to connect to, to reach this data sharing group. These names are specified in the QSGDATA parameter of the CSQ6SYSP system parameter macro (described in Using CSQ6SYSP).

By default Db2 uses the user ID of the person running the jobs as the owner of the Db2 resources. If this user ID is deleted then the resources associated with it are deleted, and so the table is deleted. Consider using a group ID to own the tables, rather than an individual user ID. You can do this by adding GROUP=groupname onto the JOB card, and specifying SET CURRENT SQLID='groupname' before any SQL statements.

IBM MQ uses the RRS Attach facility of Db2. This means that you can specify the name of a Db2 group that you want to connect to. The advantage of connecting to a Db2 group attach name (rather than a specific Db2 subsystem), is that IBM MQ can connect (or reconnect) to any available Db2 subsystem on the z/OS image that is a member of that group. There must be a Db2 subsystem that is a member of the data sharing group active on each z/OS image where you are going to run a queue-sharing IBM MQ subsystem, and RRS must be active.

Db2 storage

For most installations, the amount of Db2 storage required is about 20 or 30 cylinders on a 3390 device. However, if you want to calculate your storage requirement, the following table gives some information to help you determine how much storage Db2 requires for the IBM MQ data. The table describes the length of each Db2 row, and when each row is added to or deleted from the relevant Db2 table. Use this information together with the information about calculating the space requirements for the Db2 tables and their indexes in the *Db2 for z/OS Installation Guide*.

Table 24. Planning your Db2 storage requirements

Db2 table name	Length of row	A row is added when:	A row is deleted when:
CSQ.ADMIN_B_QSG	252 bytes	A queue-sharing group is added to the table with the ADD QSG function of the CSQ5PQSG utility.	A queue-sharing group is removed from the table with the REMOVE QSG function of the CSQ5PQSG utility. (All rows relating to this queue-sharing group are deleted automatically from all the other Db2 tables when the queue-sharing group record is deleted.)
CSQ.ADMIN_B_QMGR	Up to 3828 bytes	A queue manager is added to the table with the ADD QMGR function of the CSQ5PQSG utility.	A queue manager is removed from the table with the REMOVE QMGR function of the CSQ5PQSG utility.
CSQ.ADMIN_B_STRUCTURE	1454 bytes	The first local queue definition, specifying the QSGDISP(SHARED) attribute, that names a previously unknown structure within the queue-sharing group is defined.	The last local queue definition, specifying the QSGDISP(SHARED) attribute, that names a structure within the queue-sharing group is deleted.
CSQ.ADMIN_B_SCST	342 bytes	A shared channel is started.	A shared channel becomes inactive.

Table 24. Planning your Db2 storage requirements (continued)

Db2 table name	Length of row	A row is added when:	A row is deleted when:
CSQ.ADMIN_B_SSKT	254 bytes	A shared channel that has the NPMSPEED(NORMAL) attribute is started.	A shared channel that has the NPMSPEED(NORMAL) attribute becomes inactive.
CSQ.ADMIN_B_STRBACKUP	514 bytes	A new row is added to the CSQ.ADMIN_B_STRUCTURE table. Each entry is a dummy entry until the BACKUP CFSTRUCT command is run, which overwrites the dummy entries.	A row is deleted from the CSQ.ADMIN_B_STRUCTURE table.
CSQ.OBJ_B_AUTHINFO	3400 bytes	An authentication information object with QSGDISP(GROUP) is defined.	An authentication information object with QSGDISP(GROUP) is deleted.
CSQ.OBJ_B_QUEUE	Up to 3707 bytes	<ul style="list-style-type: none"> • A queue with the QSGDISP(GROUP) attribute is defined. • A queue with the QSGDISP(SHARED) attribute is defined. • A model queue with the DEFTYPE(SHAREDYN) attribute is opened. 	<ul style="list-style-type: none"> • A queue with the QSGDISP(GROUP) attribute is deleted. • A queue with the QSGDISP(SHARED) attribute is deleted. • A dynamic queue with the DEFTYPE(SHAREDYN) attribute is closed with the DELETE option.
CSQ.OBJ_B_NAMELIST	Up to 15127 bytes	A namelist with the QSGDISP(GROUP) attribute is defined.	A namelist with the QSGDISP(GROUP) attribute is deleted.
CSQ.OBJ_B_CHANNEL	Up to 14127 bytes	A channel with the QSGDISP(GROUP) attribute is defined.	A channel with the QSGDISP(GROUP) attribute is deleted.
CSQ.OBJ_B_STGCLASS	Up to 2865 bytes	A storage class with the QSGDISP(GROUP) attribute is defined.	A storage class with the QSGDISP(GROUP) attribute class is deleted.
CSQ.OBJ_B_PROCESS	Up to 3347 bytes	A process with the QSGDISP(GROUP) attribute is defined.	A process with the QSGDISP(GROUP) attribute is deleted.
CSQ.OBJ_B_TOPIC	Up to 14520 bytes	A topic object with QSGDISP(GROUP) attribute is defined.	A topic object with QSGDISP(GROUP) attribute is deleted.
CSQ.EXTEND_B_QMGR	Less than 430 bytes	A queue manager is added to the table with the ADD QMGR function of the CSQ5PQSG utility.	A queue manager is removed from the table with the REMOVE QMGR function of the CSQ5PQSG utility.
CSQ.ADMIN_B_MESSAGES	87 bytes	For large message PUT (1 per BLOB).	For large message GET (1 per BLOB).
CSQ.ADMIN_MSGS_BAUX1 CSQ.ADMIN_MSGS_BAUX2 CSQ.ADMIN_MSGS_BAUX3 CSQ.ADMIN_MSGS_BAUX4		These 4 tables contain message payload for large messages added into one of these 4 tables for each BLOB of the message. BLOBS are up to 511 KB in length, so if the message size is > 711 KB, there will be multiple BLOBs for this message.	

The use of large numbers of shared queue messages of size greater than 63 KB can have significant performance implications on your IBM MQ system. For more information, see SupportPac MP16, Capacity Planning and Tuning for IBM MQ for z/OS, at: Business Integration - IBM MQ SupportPacs.

Planning your logging environment

Use this topic to plan the number, size and placement of the logs, and log archives used by IBM MQ.

Logs are used to:

- Write recovery information about persistent messages
- Record information about units of work using persistent messages
- Record information about changes to objects, such as define queue
- Backup CF structures

and for other internal information.

The IBM MQ logging environment is established using the system parameter macros to specify options, such as: whether to have single or dual active logs, what media to use for the archive log volumes, and how many log buffers to have.

These macros are described in Task 14: Create the bootstrap and log data sets and Task 17: Tailor your system parameter module.

Note: If you are using queue-sharing groups, ensure that you define the bootstrap and log data sets with SHAREOPTIONS(2 3).

This section contains information about the following topics:

Log data set definitions

z/OS

Use this topic to decide on the most appropriate configuration for your log data sets.

This topic contains information to help you answer the following questions:

- Should your installation use single or dual logging?
- How many active log data sets do you need?
- “How large should the active logs be?” on page 158
- Active log placement

Should your installation use single or dual logging?

In general you should use dual logging for production, to minimize the risk of losing data. If you want your test system to reflect production, both should use dual logging, otherwise your test systems can use single logging.

With single logging data is written to one set of log data sets. With dual logging data is written to two sets of log data sets, so in the event of a problem with one log data set, such as the data set being accidentally deleted, the equivalent data set in the other set of logs can be used to recover the data.

With dual logging you require twice as much DASD as with single logging.

If you are using dual logging, then also use dual BSDSs and dual archiving to ensure adequate provision for data recovery.

Dual active logging adds a small performance cost.

Attention: Always use dual logging and dual BSDSs rather than dual writing to DASD (mirroring). If a mirrored data set is accidentally deleted, both copies are lost.

If you use persistent messages, single logging can increase maximum capacity by 10-30% and can also improve response times.

Single logging uses 2 - 310 active log data sets, whereas dual logging uses 4 - 62 active log data sets to provide the same number of active logs. Thus single logging reduces the amount of data logged, which might be important if your installation is I/O constrained.

How many active log data sets do you need?

The number of logs depends on the activities of your queue manager. For a test system with low throughput, three active log data sets might be suitable. For a high throughput production system you might want the maximum number of logs available, so, if there is a problem with offloading logs you have more time to resolve the problems.

You must have at least three active log data sets, but it is preferable to define more. For example, if the time taken to fill a log is likely to approach the time taken to archive a log during peak load, define more logs.

You should also define more logs to offset possible delays in log archiving. If you use archive logs on tape, allow for the time required to mount the tape.

Consider having enough active log space to keep a day's worth of data, in case the system is unable to archive because of lack of DASD or because it cannot write to tape.

It is possible to dynamically define new active log data sets as a way of minimizing the effect of archive delays or problems. New data sets can be brought online rapidly, using the DEFINE LOG command to avoid queue manager 'stall' due to lack of space in the active log.

If you want to define more than 31 active log data sets, you must apply APAR PI46853 and configure your logging environment to use a version 2 format BSDS. Once a version 2 format BSDS is in use, up to 310 active log data sets can be defined for each log copy ring. See "Planning to increase the maximum addressable log range" on page 163 for information on how you convert to a version 2 format BSDS.

You can tell whether your queue manager is using a version 2 or higher BSDS, either by running the print log map utility (CSQJU004), or from the CSQJ034I message issued during queue manager initialization. An end of log RBA range of FFFFFFFFFFFFFFFF, in the CSQJ034I message, indicates that a version 2, or higher, format BSDS is in use.

When a queue manager is using a version 2, or higher, format BSDS it is possible to use the `../com.ibm.mq.ref.adm.doc/q085640_dita` command to dynamically add more than 31 active log data sets to a log copy ring.

How large should the active logs be?

On IBM MQ Version 8.0, the maximum supported active log size, when archiving to disk, is 4 GB. In previous releases of the product, the maximum supported active log size when archiving to disk is 3 GB.

When archiving to tape the maximum active log size is 4 GB.

You should create active logs of at least 1 GB in size for production and test systems.

Important: You need to be careful when allocating data sets, because IDCAMS rounds up the size you allocate.

To allocate a 3 GB log specify one of the following options:

- Cylinders(4369)
- Megabytes(3071)
- TRACKS(65535)
- RECORD(786420)

Any one of these allocates 2.99995 GB.

To allocate a 4GB log specify one of the following options:

- Cylinders(5825)
- Megabytes(4095)
- TRACKS(87375)
- RECORD(1048500)

Any one of these allocates 3.9997 GB.

When using striped data sets, where the data set is spread across multiple volumes, the specified size value is allocated on each DASD volume used for striping. So, if you want to use 4 GB logs and four volumes for striping, you should specify:

- CYLinders(1456)
- Megabytes(1023)

Setting these attributes allocates $4 * 1456 = 5824$ Cylinders or $4 * 1023 = 4092$ Megabytes.

Note: Striping is supported when using extended format data sets. This is usually set by the storage manager.

Active log placement

For performance reasons you should consider striping your active log data sets. The I/O is spread across multiple volumes and reduces the I/O response times, leading to higher throughput. See the preceding text for information about allocating the size of the active logs when using striping.

You should review the I/O statistics using reports from RMF or a similar product., Perform the review of these statistics monthly (or more frequently) for the IBM MQ data sets, to ensure there are no delays due to the location of the data sets.

In some situations, there can be much IBM MQ page set I/O, and this can impact the IBM MQ log performance if they are located on the same DASD.

If you use dual logging, ensure that each set of active and archive logs is kept apart. For example, allocate them on separate DASD subsystems, or on different devices.

This reduces the risk of them both being lost if one of the volumes is corrupted or destroyed. If both copies of the log are lost, the probability of data loss is high.

When you create a new active log data, set you should preformat it using CSQJUFMT. If the log is not preformatted, the queue manager formats the log the first time it is used, which impacts the performance.

With older DASD with large spinning disks, you had to be careful which volumes were used to get the best performance.

With modern DASD, where data is spread over many PC sized disks, you do not need to worry so much about which volumes are used.

Your storage manager should be checking the enterprise DASD to review and resolve any performance problems. For availability, you might want to use one set of logs on one DASD subsystem, and the dual logs on a different DASD subsystem.

Planning your log archive storage

Use this topic to understand the different ways of maintaining your archive log data sets.

You can place archive log data sets on standard-label tapes, or DASD, and you can manage them by data facility hierarchical storage manager (DFHSM). Each z/OS logical record in an archive log data set is a VSAM control interval from the active log data set. The block size is a multiple of 4 KB.

Archive log data sets are dynamically allocated, with names chosen by IBM MQ. The data set name prefix, block size, unit name, and DASD sizes needed for such allocations are specified in the system parameter module. You can also choose, at installation time, to have IBM MQ add a date and time to the archive log data set name.

It is not possible to specify with IBM MQ, specific volumes for new archive logs, but you can use Storage Management routines to manage this. If allocation errors occur, offloading is postponed until the next time offloading is triggered.

If you specify dual archive logs at installation time, each log control interval retrieved from the active log is written to two archive log data sets. The log records that are contained in the pair of archive log data sets are identical, but the end-of-volume points are not synchronized for multivolume data sets.

Should your archive logs reside on tape or DASD?

When deciding whether to use tape or DASD for your archive logs, there are a number of factors that you should consider:

- Review your operating procedures before deciding about tape or disk. For example, if you choose to archive to tape, there must be enough tape drive when they are required. After a disaster, all subsystems might want tape drives and you might not have as many free tape drives as you expect.
- During recovery, archive logs on tape are available as soon as the tape is mounted. If DASD archives have been used, and the data sets migrated to tape using hierarchical storage manager (HSM), there is a delay while HSM recalls each data set to disk. You can recall the data sets before the archive log is used. However, it is not always possible to predict the correct order in which they are required.
- When using archive logs on DASD, if many logs are required (which might be the case when recovering a page set after restoring from a backup) you might require a significant quantity of DASD to hold all the archive logs.
- In a low-usage system or test system, it might be more convenient to have archive logs on DASD to eliminate the need for tape mounts.
- Both issuing a RECOVER CFSTRUCT command, and backing out a persistent unit of work, result in the log being read backwards. Tape drives with hardware compression perform badly on operations that read backwards. Plan sufficient log data on DASD to avoid reading backwards from tape.

Archiving to DASD offers faster recoverability but is more expensive than archiving to tape. If you use dual logging, you can specify that the primary copy of the archive log go to DASD and the secondary copy go to tape. This increases recovery speed without using as much DASD, and you can use the tape as a backup.

Archiving to tape

If you choose to archive to a tape device, IBM MQ can extend to a maximum of 20 volumes.

If you are considering changing the size of the active log data set so that the set fits on one tape volume, note that a copy of the BSDS is placed on the same tape volume as the copy of the active log data set. Adjust the size of the active log data set downward to offset the space required for the BSDS on the tape volume.

If you use dual archive logs on tape, it is typical for one copy to be held locally, and the other copy to be held off-site for use in disaster recovery.

Archiving to DASD volumes

IBM MQ requires that you catalog all archive log data sets allocated on non-tape devices (DASD). If you choose to archive to DASD, the CATALOG parameter of the CSQ6ARVP macro must be YES. If this parameter is NO, and you decide to place archive log data sets on DASD, you receive message CSQJ072E each time an archive log data set is allocated, although IBM MQ still catalogs the data set.

If the archive log data set is held on DASD, the archive log data sets can extend to another volume; multivolume is supported.

If you choose to use DASD, make sure that the primary space allocation (both quantity and block size) is large enough to contain either the data coming from the active log data set, or that from the corresponding BSDS, whichever is the larger of the two.

This minimizes the possibility of unwanted z/OS X'B37' or X'E37' abend codes during the offload process. The primary space allocation is set with the PRIQTY (primary quantity) parameter of the CSQ6ARVP macro.

From IBM MQ Version 8.0, archive log data sets can exist on large or extended-format sequential data sets. SMS ACS routines can now use DSNTYPE(LARGE) or DSNTYPE(EXT). These were not supported before Version 8.0.

IBM MQ supports allocation of archive logs as extended format data sets. When extended format is used, the maximum archive log size is increased from 65535 tracks to the maximum active log size of 4GB. Archive logs are eligible for allocation in the extended addressing space (EAS) of extended address volumes (EAV).

Where the required hardware and software levels are available, allocating archive logs to a data class defined with COMPACTION using zEDC might reduce the disk storage required to hold archive logs. For more information, see IBM MQ for z/OS: Reducing storage occupancy with IBM zEnterprise Data Compression (zEDC).

Refer to Using the zEnterprise Data Compression (zEDC) enhancements for details on hardware and software levels, as well as example RACF profile changes.

The z/OS dataset encryption feature can be applied to archive logs for queue managers running at IBM MQ Version 8.0 or later. These archive logs must be allocated through Automatic Class Selection (ACS) routines to a data class defined with EXTENDED attributes, and a data set key label that ensures the data is AES encrypted.

Using SMS with archive log data sets

If you have MVS™/DFP storage management subsystem (DFSMS) installed, you can write an Automatic Class Selection (ACS) user-exit filter for your archive log data sets, which helps you convert them for the SMS environment.

Such a filter, for example, can route your output to a DASD data set, which DFSMS can manage. You must exercise caution if you use an ACS filter in this manner. Because SMS requires DASD data sets to be cataloged, you must make sure the CATALOG DATA field of the CSQ6ARVP macro contains YES. If it does not, message CSQJ072E is returned; however, the data set is still cataloged by IBM MQ.

For more information about ACS filters, see DFP Storage Administration Reference.

How long do I need to keep archive logs

Use the information in this section to help you plan your backup strategy.

You specify how long archive logs are kept in days , using the ARCRETN parameter in USING CSQ6ARVP or the SET SYSTEM command. After this period the data sets can be deleted by z/OS.

You can manually delete archive log data sets when they are no longer needed.

- The queue manager might need the archive logs for recovery.
The queue manager can only keep the most recent 1000 archives in the BSDS, When the archive logs are not in the BSDS they cannot be used for recovery, and are only of use for audit, analysis, or replay type purposes.
- You might want to keep the archive logs so that you can extract information from the logs. For example, extracting messages from the log, and reviewing which user ID put or got the message.

The BSDS contains information on logs and other recovery information. This dataset is a fixed size. When the number of archive logs reaches the value of MAXARCH in CSQ6LOGP, or when the BSDS fills up, the oldest archive log information is overwritten.

There are utilities to remove archive log entries from the BSDS, but in general, the BSDS wraps and overlays the oldest archive log record.

When is an archive log needed

You need to backup your page sets regularly. The frequency of backups determines which archive logs are needed in the event of losing a page set.

You need to backup your CF structures regularly. The frequency of backups determines which archive logs are needed in the event of losing data in the CF structure.

The archive log might be needed for recovery. The following information explains when the archive log might be needed, where there are problems with different IBM MQ resources.

Loss of Page set 0

You must recover your system from your backup and restart the queue manager.

You need the logs from when the backup was taken, plus up to three active logs.

Loss of any other page set

You must recover your system from your backup and restart the queue manager.

You need the logs from when the backup was taken, plus up to three active logs.

All LPARS lose connectivity to a structure, or the structure is unavailable

Use the RECOVER CFSTRUCT command to read from the last CF backup on the logs.

If you have been doing frequent backups of the CF, the data should be in active logs.

You should not need archive logs.

Administration structure rebuild

If you need to rebuild the administration structure, the information is read from the last checkpoint of the log for each queue manager.

If a queue manager is not active, another queue manager reads the log.

You should not need archive logs.

Loss of an SMDS dataset

If you lose an SMDS dataset, or the dataset gets corrupted, the dataset becomes unusable and the status for it is set to FAILED. The CF structure is unchanged.

In order to restore the SMDS dataset, you need to:

1. Redefine the SMDS dataset, and
2. Fail and then recover the CF structure.

Issuing the RECOVER CFSTRUCT command twice achieves this process.

Issuing the command the first time sets the structure state to failed; issuing the command a second time does the actual recovery.

Note: All non persistent messages on the CF structure will be lost; all persistent messages will be restored.

You will need the logs from the time the BACKUP CFSTRUCT command was issued, so this might require archive logs.

If all LPARs lose connectivity to the structure, the structure is recreated, possibly in an alternative CF. Note that your structure CFRM PREFLIST attribute must contain multiple CFs.

Note: All non persistent messages will be lost; all persistent messages will be recreated by:

1. Reading the log for the last CF backup
2. Reading the logs from all queue managers that have used the structure, and
3. Merging updates since the backup

You require the logs from all queue managers that have accessed the structure since the last backup (back to the time when the backup was taken) plus the structure backup itself in the log of the queue manager that took the backup.

BSDS

Do you need single or dual BSDS?

If you are using dual active logs you should use dual BSDS.

How big does the BSDS need to be?

The BSDS does not need to be very large, and a primary and secondary of one cylinder should be sufficient.

Planning to increase the maximum addressable log range

You can increase the maximum addressable log range by configuring your queue manager to use a larger log relative byte address (RBA).

For an overview of the change to the log RBA for IBM MQ Version 8.0 , see Larger log Relative Byte Address.

If the queue manager is not in a queue-sharing group, you can upgrade the queue manager to IBM MQ Version 8.0 , enable Version 8.0 new functions, and convert it to use 8 byte log RBA values at any time. Once a queue manager has been converted to use 8 byte log RBA values, it is not possible to revert it to COMPAT mode.

For queue managers in a queue-sharing group, you can upgrade each queue manager in turn to IBM MQ Version 8.0 and enable Version 8.0 new function. Once all the queue managers in the group are at IBM MQ Version 8.0 new function mode, you can change each queue manager in turn to use 8 byte log RBA values. It is not essential to change all the queue managers at the same time.

When a queue manager in a queue-sharing group has been converted to use 8 byte log RBA values, other queue managers in the queue-sharing group can use the logs of the converted queue manager, even though they have not yet been converted to use 8 byte log RBA values. This is useful, for example, for peer recovery.

Note: A queue manager that has been converted to use 8 byte log RBA values can read logs that have data written with 6 byte or 8 byte log RBA values. Therefore, its active logs and archive logs can be used to recover page sets and coupling facility (CF) structures.

Undoing the change

The change cannot be backed out.

How long does it take?

The change requires a queue manager restart. Stop the queue manager, run the CSQJUCNV utility against the bootstrap data set (BSDS), or data sets, to create new data sets, rename these bootstrap data sets, and restart the queue manager.

What impact does this have?

- With 8 byte log RBA in use, every write of data to the log data sets has additional bytes. Therefore, for a workload consisting of persistent messages there is a small increase in the amount of data written to the logs.
- Data written to a page set, or coupling facility (CF) structure, is not affected.

Related information:

Implementing the larger log Relative Byte Address

Planning for IBM MQ Managed File Transfer

Use this topic as guidance on how you need to set up your system to run IBM MQ Managed File Transfer .

Common configurations

There are three common IBM MQ Managed File Transfer (MFT) configurations:

1. A single queue manager with one or more agents using local connections. This might be used to put the contents of a data set into IBM MQ queues.
2. A single queue manager with an MFT client on a distributed machine using client bindings.
3. Two queue managers connected by channels, and one or more agents on each machine. These agents can be client or local bindings.

MFT can use multiple queue managers:

- One or more queue managers to transfer the data.
- A commands queue manager that issues requests. For example, a request to start a transfer is sent to this queue manager, and the associated commands are routed to the MFT agents.
- A coordination queue manager that manages the work.

Notes:

1. You can use the same queue manager for transferring data, commands and coordination.
2. This setup, although the simplest, might not be the most efficient because all the workload is on one queue manager.

If you have an existing IBM MQ Managed File Transfer configuration, your command and coordination queue manager might already exist.

If you do not have an existing IBM MQ Managed File Transfer configuration, you can use one queue manager for transferring data, commands, and coordination. Note that even if you do this, it is possible to set up multiple configurations on the same machine.

If you are using multiple queue managers you need to set up channels between the queue managers. You can either do this by using clustering or by using point-to-point connections. IBM MQ Managed File Transfer status and activity can be logged, and can be stored in either a Db2 or Oracle database.

IBM MQ Managed File Transfer is written in Java, with some shell scripts and JCL to configure and operate the program.

Important: You must be familiar with UNIX System Services (USS) in order to configure IBM MQ Managed File Transfer. For example:

- The file directory structure, with names such as /u/userID/myfile.txt
- USS commands, for example:
 - cd (change directory)
 - ls (list)
 - chmod (change the file permissions)
 - chown (change file ownership or groups which can access the file or directory)

You require the following products in USS to be able to configure and run MFT:

1. Java, for example, in directory /java/java71_bit64_GA/J7.1_64/
2. IBM MQ Version 8.0 , for example, in directory /mqm/V8R0M0
3. If you want to use Db2 for status and history, you need to install Db2 JDBC libraries, for example, in directory /db2/db2v10/jdbc/libs.

Product registration

At startup IBM MQ Managed File Transfer checks the registration in sys1.parmlib concatenation. The following code is an example of how you register MFT:

```
PRODUCT OWNER('IBM CORP')
NAME('WS MQ FILE TRANS')
ID(5655-MFT)
VERSION(*) RELEASE(*) MOD(*)
FEATURENAME('WS MQ FILE TRANS')
STATE(ENABLED)
```

Disk space

You will need 100 MB of DASD for PDSEs and a minimum of 50 MB in USS. If you used trace to diagnose problems, you need additional disk space in USS, for example 50 MB.

Security

You need to identify which user IDs are going to be used for MFT configuration and for MFT operation.

You need to identify the files or queues you transfer, and which user IDs are going to be submitting transfer requests to MFT.

When you customize the agents and logger, you specify the group of users that is allowed to run MFT services, or do MFT administration.

You should set up this group before you start customizing MFT. As MFT uses IBM MQ queues, if you have security enabled in the queue manager, MFT requires access to the following resources

Table 25.

Name	Access required
QUEUE.SYSTEM.FTE.EVENT.agent_name	Update
QUEUE.SYSTEM.FTE.COMMAND.agent_name	Update
CONTEXT.SYSTEM.FTE.COMMAND.agent_name	Update
QUEUE.SYSTEM.FTE.STATE.agent_name	Update
QUEUE.SYSTEM.FTE.DATA.agent_name	Update
QUEUE.SYSTEM.FTE.REPLY.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHAGT1.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHTRN1.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHOPS1.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHSCH1.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHMON1.agent_name	Update
QUEUE.SYSTEM.FTE.AUTHADM1.agent_name	Update

Table 26.

Name	Access required
SYSTEM.FTE.AUTHAGT1.agent_name	Update
SYSTEM.FTE.AUTHTRN1.agent_name	Update
SYSTEM.FTE.AUTHOPS1.agent_name	Update
SYSTEM.FTE.AUTHSCH1.agent_name	Update
SYSTEM.FTE.AUTHMON1.agent_name	Update

You can use user sandboxing to determine which parts of the file system the user who requests the transfer can access.

To enable user sandboxing, add the `userSandboxes=true` statement to the `agent.properties` file for the agent that you want to restrict, and add appropriate values to the `MQ_DATA_PATH/mqft/config/coordination_qmgr_name/agents/agent_name/UserSandboxes.xml` file.

See [Working with user sandboxes](#) for further information.

This user ID is configured in `UserSandboxes.xml` files.

This XML file has information like user ID, or user ID* and a list of resource that can be used (included), or cannot be used (excluded). You need to define specific user IDs that can access which resources: for example:

Table 27.

User ID	Access	Include or Exclude	Resource
Admin*	Read	Include	/home/user/**
Admin*	Read	Exclude	/home/user/private/**
Sysprog	Read	Include	/home/user/**
Admin*	Read	Include	Application.reply.queue

Notes:

1. If type=queue is specified, the resource is either a queue name, or queue@qmgr.
2. If the resource begins with //, the resource is a dataset; otherwise the resource is a file in USS.
3. The user ID is the user ID from the MQMD structure, so this might not reflect the user ID that actually puts the message.
4. For requests on the local queue manager you can use MQADMIN CONTEXT.* to limit which users can set this value.
5. For requests coming in over a remote queue manager, you have to assume that the distributed queue managers have security enabled to prevent unauthorized setting of the user ID in the MQMD structure.
6. A user ID of SYSPROG1 on a Linux machine, is the same user ID SYSPROG1 for the security checking on z/OS.

How many agents do I need?

The agents do the work in transferring data, and when you make a request to transfer data you specify the name of an agent.

By default an agent can process 25 send and 25 receive requests concurrently. You can configure these processes. See IBM MQ Managed File Transfer configuration options on z/OS for more information.

If the agent is busy then work is queued. The time taken to process a request depends on multiple factors, for example, the amount of data to be sent, the network bandwidth, and the delay on the network.

You might want to have multiple agents to process work in parallel.

You can also control which resources an agent can access, so you might want some agents to work with a limited subset of data.

If you want to process requests with different priority you can use multiple agents and use workload manager to set the priority of the jobs.

Running the agents

Typically the agents are long running processes. The processes can be submitted as jobs that run in batch, or as started tasks.

Planning for channel initiator SMF data

You need to plan the implementation of SMF data for the channel initiator (CHINIT).

The CHINIT produces two types of record:

- Statistics data with information about the CHINIT and the tasks within it.
- Channel accounting data with information similar to the DIS CHSTATUS command.

You start collecting statistics data using:

```
/CPF START TRACE(S) class(4)
```

and stop it using

```
/CPF STOP TRACE(S) class(4)
```

You start collecting accounting data using:

```
/CPF START TRACE(A) class(4)
```

and stop it using

```
/CPF STOP TRACE(A) class(4)
```

The SMF records are produced when:

- The time interval in the STATIME ZPARM parameter has elapsed, or if STATIME is zero, on the SMF broadcast. The request to collect SMF data for the CHINIT and the queue manager are synchronized.
- A STOP TRACE(A) CLASS(4) or STOP TRACE(S) CLASS(4) command is issued, or
- When the CHINIT is shut down. At this point any SMF is written out.

The statistics SMF data normally fits into one SMF record, however, multiple SMF records might be created if a large number of tasks are in use.

Accounting data is gathered for each channel for which it is enabled (using the STATCHL attribute) and normally fits into one SMF record. However, multiple SMF records might be created if a large number of channels are active.

If a channel stops in the interval, accounting data is written to SMF the next time the SMF processing runs. If a client connects, does some work and disconnects, then reconnects and disconnects, there are two sets of channel accounting data produced.

You can control which channels have information written to SMF:

1. By using the STATCHL option on the channel and queue manager.
2. For client channels, note that, you must set STATCHL at the queue manager level.
3. For automatically defined cluster sender channels, you must set STATACLS.

The cost of using the CHINIT SMF data is small. Typically the increase in CPU usage is under a few percent, and often within measurement error.

Before you use this function you need to work with your z/OS systems programmer to ensure that SMF has the capacity for the additional records, and that they change their processes for extracting SMF records to include the new SMF data.

For CHINIT statistics the SMF record type is 115 and sub-type 231.

For CHINIT accounting the SMF record type is 116 and sub-type 10.

You can write your own programs to process this data, or use the SupportPac MP1B that contains a program, MQSMF, for printing the data, and creating data in Comma Separated Values (CSV) suitable for importing into a spread sheet.

If you are experiencing issues with capturing channel initiator SMF data, see [Dealing with issues when capturing SMF data for the channel initiator \(CHINIT\)](#) for further information.

Related information:

Interpreting IBM MQ performance statistics
Troubleshooting channel accounting data

Planning for backup and recovery

Developing backup and recovery procedures at your site is vital to avoid costly and time-consuming losses of data. IBM MQ provides means for recovering both queues and messages to their current state after a system failure.

This topic contains the following sections:

- “Recovery procedures”
- “Tips for backup and recovery” on page 170
- “Recovering page sets” on page 172
- “Recovering CF structures” on page 173
- “Achieving specific recovery targets” on page 174
- “Backup considerations for other products” on page 175
- “Recovery and CICS” on page 175
- “Recovery and IMS” on page 176
- “Preparing for recovery on an alternative site” on page 176
- “Example of queue manager backup activity” on page 176

Recovery procedures

Develop the following procedures for IBM MQ:

- Creating a point of recovery.
- Backing up page sets.
- Backing up CF structures.
- Recovering page sets.
- Recovering from out-of-space conditions (IBM MQ logs and page sets).
- Recovering CF structures.

See Administering IBM MQ for z/OS for information about these.

Become familiar with the procedures used at your site for the following:

- Recovering from a hardware or power failure.
- Recovering from a z/OS component failure.
- Recovering from a site interruption, using off-site recovery.

Tips for backup and recovery

Use this topic to understand some backup and recovery tasks.

The queue manager restart process recovers your data to a consistent state by applying log information to the page sets. If your page sets are damaged or unavailable, you can resolve the problem using your backup copies of your page sets (if all the logs are available). If your log data sets are damaged or unavailable, it might not be possible to recover completely.

Consider the following points:

- Periodically take backup copies
- Do not discard archive logs you might need
- Do not change the DDname to page set association

Periodically take backup copies

A *point of recovery* is the term used to describe a set of backup copies of IBM MQ page sets and the corresponding log data sets required to recover these page sets. These backup copies provide a potential restart point in the event of page set loss (for example, page set I/O error). If you restart the queue manager using these backup copies, the data in IBM MQ is consistent up to the point that these copies were taken. Provided that all logs are available from this point, IBM MQ can be recovered to the point of failure.

The more recent your backup copies, the quicker IBM MQ can recover the data in the page sets. The recovery of the page sets is dependent on all the necessary log data sets being available.

In planning for recovery, you need to determine how often to take backup copies and how many complete backup cycles to keep. These values tell you how long you must keep your log data sets and backup copies of page sets for IBM MQ recovery.

When deciding how often to take backup copies, consider the time needed to recover a page set. The time needed is determined by the following:

- The amount of log to traverse.
- The time it takes an operator to mount and remove archive tape volumes.
- The time it takes to read the part of the log needed for recovery.
- The time needed to reprocess changed pages.
- The storage medium used for the backup copies.
- The method used to make and restore backup copies.

In general, the more frequently you make backup copies, the less time recovery takes, but the more time is spent making copies.

For each queue manager, you should take backup copies of the following:

- The archive log data sets
- The BSDS copies created at the time of the archive
- The page sets
- Your object definitions
- Your CF structures

To reduce the risk of your backup copies being lost or damaged, consider:

- Storing the backup copies on different storage volumes to the original copies.
- Storing the backup copies at a different site to the original copies.

- Making at least two copies of each backup of your page sets and, if you are using single logging or a single BSDS, two copies of your archive logs and BSDS. If you are using dual logging or BSDS, make a single copy of both archive logs or BSDS.

Before moving IBM MQ to a production environment, fully test and document your backup procedures.

Backing up your object definitions

Create backup copies of your object definitions. To do this, use the MAKEDEF feature of the COMMAND function of the utility program (described in Using the COMMAND function of CSQUTIL).

You should do this whenever you take backup copies of your queue manager data sets, and keep the most current version.

Backing up your coupling facility structures

If you have set up any queue-sharing groups, even if you are not using them, you must take periodic backups of your CF structures. To do this, use the IBM MQ BACKUP CFSTRUCT command. You can use this command only on CF structures that are defined with the RECOVER(YES) attribute. If any CF entries for persistent shared messages refer to offloaded message data stored in a shared message data set (SMDS) or Db2, the offloaded data is retrieved and backed up together with the CF entries. Shared message data sets should not be backed up separately.

It is recommended that you take a backup of all your CF structures about every hour, to minimize the time it takes to restore a CF structure.

You could perform all your CF structure backups on a single queue manager, which has the advantage of limiting the increase in log use to a single queue manager. Alternatively, you could perform backups on all the queue managers in the queue-sharing group, which has the advantage of spreading the workload across the queue-sharing group. Whichever strategy you use, IBM MQ can locate the backup and perform a RECOVER CFSTRUCT from any queue manager in the queue-sharing group. The logs of all the queue managers in the queue-sharing group need to be accessed to recover the CF structure.

Backing up your message security policies

If you are using IBM MQ Advanced Message Security to create a backup of your message security policies, create a backup using the message security policy utility (CSQ0UTIL) to run **dspmqspl** with the -export parameter, then save the policy definitions that are output to the EXPORT DD.

You should create a backup of your message security policies whenever you take backup copies of your queue manager data sets, and keep the most current version.

Do not discard archive logs you might need

IBM MQ might need to use archive logs during restart. You must keep sufficient archive logs so that the system can be fully restored. IBM MQ might use an archive log to recover a page set from a restored backup copy. If you have discarded that archive log, IBM MQ cannot restore the page set to its current state. When and how you discard archive logs is described in Discarding archive log data sets.

You can use the /cpf DIS USAGE TYPE(ALL) command to display the log RBA, and log range sequence number (LRSN) that you need to recover your queue manager's page sets and the queue-sharing group's structures. You should then use the print log map utility (CSQJU004) to print bootstrap data set (BSDS) information for the queue manager to locate the logs containing the log RBA.

For structures, you need to run the CSQJU004 utility on each queue manager in the queue-sharing group to locate the logs containing the LRSN. You need these logs and any later logs to be able to recover the page sets and structures.

Do not change the DDname to page set association

IBM MQ associates page set number 00 with DDname CSQP0000, page set number 01 with DDname CSQP0001, and so on, up to CSQP0099. IBM MQ writes recovery log records for a page set based on the DDname that the page set is associated with. For this reason, you must not move page sets that have already been associated with a PSID DDname.

Recovering page sets

Use this topic to understand the factors involved when recovering pages sets, and how to minimize restart times.

A key factor in recovery strategy concerns the time for which you can tolerate a queue manager outage. The total outage time might include the time taken to recover a page set from a backup, or to restart the queue manager after an abnormal termination. Factors affecting restart time include how frequently you back up your page sets, and how much data is written to the log between checkpoints.

To minimize the restart time after an abnormal termination, keep units of work short so that, at most, two active logs are used when the system restarts. For example, if you are designing an IBM MQ application, avoid placing an **MQGET** call that has a long wait interval between the first in-syncpoint MQI call and the commit point because this might result in a unit of work that has a long duration. Another common cause of long units of work is batch intervals of more than 5 minutes for the channel initiator.

You can use the **DISPLAY THREAD** command to display the RBA of units of work and to help resolve the old ones.

How often must you back up a page set?

Frequent page set backup is essential if a reasonably short recovery time is required. This applies even when a page set is very small or there is a small amount of activity on queues in that page set.

If you use persistent messages in a page set, the backup frequency should be in hours rather than days. This is also the case for page set zero.

To calculate an approximate backup frequency, start by determining the target total recovery time. This consists of the following:

1. The time taken to react to the problem.
2. The time taken to restore the page set backup copy.
If you use SnapShot backup/restore, the time taken to perform this task is a few seconds. For information about SnapShot, see the *DFSMSdss Storage Administration Guide*.
3. The time the queue manager requires to restart, including the additional time needed to recover the page set.

This depends most significantly on the amount of log data that must be read from active and archive logs since that page set was last backed up. All such log data must be read, in addition to that directly associated with the damaged page set.

Note: When using *fuzzy backup* (where a snapshot is taken of the logs and page sets while a unit of work is active), it might be necessary to read up to three additional checkpoints, and this might result in the need to read one or more additional logs.

When deciding on how long to allow for the recovery of the page set, the factors that you need to consider are:

- The rate at which data is written to the active logs during normal processing depends on how messages arrive in your system, in addition to the message rate.

Messages received or sent over a channel result in more data logging than messages generated and retrieved locally.

- The rate at which data can be read from the archive and active logs.

When reading the logs, the achievable data rate depends on the devices used and the total load on your particular DASD subsystem.

With most tape units, it is possible to achieve higher data rates for archived logs with a large block size. However, if an archive log is required for recovery, all the data on the active logs must be read also.

Recovering CF structures

Use this topic to understand the recovery process for CF structures.

At least one queue manager in the queue-sharing group must be active to process a RECOVER CFSTRUCT command. CF structure recovery does not affect queue manager restart time, because recovery is performed by an already active queue manager.

The recovery process consists of two logical steps that are managed by the RECOVER CFSTRUCT command:

1. Locating and restoring the backup.
2. Merging all the logged updates to persistent messages that are held on the CF structure from the logs of all the queue managers in the queue-sharing group that have used the CF structure, and applying the changes to the backup.

The second step is likely to take much longer because a lot of log data might need to be read. You can reduce the time taken if you take frequent backups, or if you recover multiple CF structures at the same time, or both.

The queue manager performing the recovery locates the relevant backups on all the other queue managers' logs using the data in Db2 and the bootstrap data sets. The queue manager replays these backups in the correct time sequence across the queue sharing group, from just before the last backup through to the point of failure.

The time it takes to recover a CF structure depends on the amount of recovery log data that must be replayed, which in turn depends on the frequency of the backups. In the worst case, it takes as long to read a queue manager's log as it did to write it. So if, for example, you have a queue-sharing group containing six queue managers, an hour's worth of log activity could take six hours to replay. In general it takes less time than this, because reading can be done in bulk, and because the different queue manager's logs can be read in parallel. As a starting point, we recommend that you back up your CF structures every hour.

All queue managers can continue working with non-shared queues and queues in other CF structures while there is a failed CF structure. If the administration structure has also failed, at least one of the queue managers in the queue-sharing group must be started before you can issue the RECOVER CFSTRUCT command.

Backing up CF structures can require considerable log writing capacity, and can therefore impose a large load on the queue manager doing the backup. Choose a lightly loaded queue manager for doing backups; for busy systems, add an additional queue manager to the queue-sharing group and dedicate it exclusively for doing backups.

Achieving specific recovery targets

Use this topic for guidance on how you can achieve specific recovery target times by adjusting backup frequency.

If you have specific recovery targets to achieve, for example, completion of the queue manager recovery and restart processing in addition to the normal startup time within *xx* seconds, you can use the following calculation to estimate your backup frequency (in hours):

Formula (A)	
Backup frequency (in hours)	$= \frac{\text{Required restart time (in secs)} * \text{System recovery log read rate (in MB/sec)}}{\text{Application log write rate (in MB/hour)}}$

Note: The examples given next are intended to highlight the need to back up your page sets frequently. The calculations assume that most log activity is derived from a large number of persistent messages. However, there are situations where the amount of log activity is not easily calculated. For example, in a queue-sharing group environment, a unit of work in which shared queues are updated in addition to other resources might result in UOW records being written to the IBM MQ log. For this reason, the 'Application log write rate' in Formula (A) can be derived accurately only from the observed rate at which the IBM MQ logs fill.

For example, consider a system in which IBM MQ MQI clients generate a total load of 100 persistent messages a second. In this case, all messages are generated locally.

If each message is of user length 1 KB, the amount of data logged each hour is approximately:

$100 * (1 + 1.3) \text{ KB} * 3600 = \text{approximately } 800 \text{ MB}$
where
100 = the message rate a second
(1 + 1.3) KB = the amount of data logged for each 1 KB of persistent messages

Consider an overall target recovery time of 75 minutes. If you have allowed 15 minutes to react to the problem and restore the page set backup copy, queue manager recovery and restart must then complete within 60 minutes (3600 seconds) applying formula (A). Assuming that all required log data is on RVA2-T82 DASD, which has a recovery rate of approximately 2.7 MB a second, this necessitates a page set backup frequency of at least every:

$3600 \text{ seconds} * 2.7 \text{ MB a second} / 800 \text{ MB an hour} = 12.15 \text{ hours}$

If your IBM MQ application day lasts approximately 12 hours, one backup each day is appropriate. However, if the application day lasts 24 hours, two backups each day is more appropriate.

Another example might be a production system in which all the messages are for request-reply applications (that is, a persistent message is received on a receiver channel and a persistent reply message is generated and sent down a sender channel).

In this example, the achieved batch size is one, and so there is one batch for every message. If there are 50 request replies a second, the total load is 100 persistent messages a second. If each message is 1 KB in length, the amount of data logged each hour is approximately:

$50((2 * (1+1.3) \text{ KB}) + 1.4 \text{ KB} + 2.5 \text{ KB}) * 3600 = \text{approximately } 1500 \text{ MB}$

where:

50 = the message pair rate a second
(2 * (1 + 1.3) KB) = the amount of data logged for each message pair
1.4 KB = the overhead for each batch of messages received by each channel
2.5 KB = the overhead for each batch of messages sent by each channel

To achieve the queue manager recovery and restart within 30 minutes (1800 seconds), again assuming that all required log data is on RVA2-T82 DASD, this requires that page set backup is carried out at least every:

$1800 \text{ seconds} * 2.7 \text{ MB a second} / 1500 \text{ MB an hour} = 3.24 \text{ hours}$

Periodic review of backup frequency

Monitor your IBM MQ log usage in terms of MB an hour. Periodically perform this check and amend your page set backup frequency if necessary.

Backup considerations for other products

If you are using IBM MQ with CICS or IMS then you must also consider the implications for your backup strategy with those products. The data facility hierarchical storage manager (DFHSM) manages data storage, and can interact with the storage used by IBM MQ.

Backup and recovery with DFHSM

The data facility hierarchical storage manager (DFHSM) does automatic space-availability and data-availability management among storage devices in your system. If you use it, you need to know that it moves data to and from the IBM MQ storage automatically.

DFHSM manages your DASD space efficiently by moving data sets that have not been used recently to alternative storage. It also makes your data available for recovery by automatically copying new or changed data sets to tape or DASD backup volumes. It can delete data sets, or move them to another device. Its operations occur daily, at a specified time, and allow for keeping a data set for a predetermined period before deleting or moving it.

You can also perform all DFHSM operations manually. The *Data Facility Hierarchical Storage Manager User's Guide* explains how to use the DFHSM commands. If you use DFHSM with IBM MQ, note that DFHSM does the following:

- Uses cataloged data sets.
- Operates on page sets and logs.
- Supports VSAM data sets.

Recovery and CICS

The recovery of CICS resources is not affected by the presence of IBM MQ. CICS recognizes IBM MQ as a non-CICS resource (or external resource manager), and includes IBM MQ as a participant in any syncpoint coordination requests using the CICS resource manager interface (RMI). For more information about CICS recovery, see the *CICS Recovery and Restart Guide*. For information about the CICS resource manager interface, see the *CICS Customization Guide*.

Recovery and IMS

IMS recognizes IBM MQ as an external subsystem and as a participant in syncpoint coordination. IMS recovery for external subsystem resources is described in the *IMS Customization Guide*.

Preparing for recovery on an alternative site

If a total loss of an IBM MQ computing center, you can recover on another IBM MQ system at a recovery site.

To recover an IBM MQ system at a recovery site, you must regularly back up the page sets and the logs. As with all data recovery operations, the objectives of disaster recovery are to lose as little data, workload processing (updates), and time as possible.

At the recovery site:

- The recovery IBM MQ queue manager **must** have the same name as the lost queue manager.
- Ensure the system parameter module used on the recovery queue manager contains the same parameters as the lost queue manager.

The process for disaster recovery is described in the *Administering IBM MQ for z/OS*.

Example of queue manager backup activity

This topic shows as an example of queue manager backup activity.

When you plan your queue manager backup strategy, a key consideration is retention of the correct amount of log data. Managing the logs describes how to determine which log data sets are required, by reference to the system recovery RBA of the queue manager. IBM MQ determines the system recovery RBA using information about the following:

- Currently active units of work.
- Page set updates that have not yet been flushed from the buffer pools to disk.
- CF structure backups, and whether this queue manager's log contains information required in any recovery operation using them.

You must retain sufficient log data to be able to perform media recovery. While the system recovery RBA increases over time, the amount of log data that must be retained only decreases when subsequent backups are taken. CF structure backups are managed by IBM MQ, and so are taken into account when reporting the system recovery RBA. This means that in practice, the amount of log data that must be retained only reduces when page set backups are taken.

Figure 57 on page 177 shows an example of the backup activity on a queue manager that is a member of a queue-sharing group, how the recovery RBA varies with each backup, and how that affects the amount of log data that must be retained. In the example the queue manager uses local and shared resources: page sets, and two CF structures, STRUCTURE1 and STRUCTURE2.

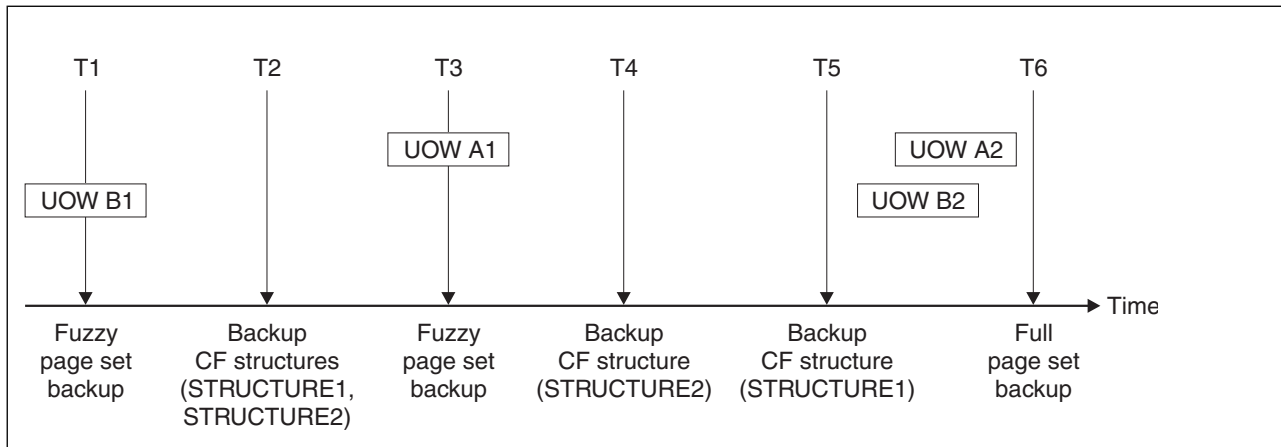


Figure 57. Example of queue manager backup activity.

Example of queue manager backup activity.

This is what happens at each point in time:

Point in time T1

A fuzzy backup is created of your page sets, as described in How to back up and recover page sets.

The system recovery RBA of the queue manager is the lowest of the following:

- The recovery RBAs of the page sets being backed up at this point.
- The lowest recovery RBA required to recover the CF application structures. This relates to the recovery of backups of STRUCTURE1 and STRUCTURE2 created earlier.
- The recovery RBA for the oldest currently active unit of work within the queue manager (UOWB1).

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the fuzzy backup process.

Point in time T2

Backups of the CF structures are created. CF structure STRUCTURE1 is backed up first, followed by STRUCTURE2.

The amount of log data that must be retained is unchanged, because the same data as determined from the system recovery RBA at T1 is still required to recover using the page set backups taken at T1.

Point in time T3

Another fuzzy backup is created.

The system recovery RBA of the queue manager is the lowest of the following:

- The recovery RBAs of the page sets being backed up at this point.
- The lowest recovery RBA required to recover CF structure STRUCTURE1, because STRUCTURE1 was backed up before STRUCTURE2.
- The recovery RBA for the oldest currently active unit of work within the queue manager (UOWA1).

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the fuzzy backup process.

You can now reduce the log data retained, as determined by this new system recovery RBA.

Point in time T4

A backup is taken of CF structure STRUCTURE2. The recovery RBA for the recovery of the oldest required CF structure backup relates to the backup of CF structure STRUCTURE1, which was backed up at time T2.

The creation of this CF structure backup has no effect on the amount of log data that must be retained.

Point in time T5

A backup is taken of CF structure STRUCTURE1. The recovery RBA for recovery of the oldest required CF structure backup now relates to recovery of CF structure STRUCTURE2, which was backed up at time T4.

The creation of this CF structure backup has no effect on amount of log data that must be retained.

Point in time T6

A full backup is taken of your page sets as described in How to back up and recover page sets.

The system recovery RBA of the queue manager is the lowest of the following:

- The recovery RBAs of the page sets being backed up at this point.
- The lowest recovery RBA required to recover the CF structures. This relates to recovery of CF structure STRUCTURE2.
- The recovery RBA for the oldest currently active unit of work within the queue manager. In this case, there are no current units of work.

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the full backup process.

Again, the log data retained can be reduced, because the system recovery RBA associated with the full backup is more recent.

Planning your z/OS UNIX or UNIX System Services environment

Certain processes within the IBM MQ queue manager (MSTR) and the channel initiator (CHIN) use z/OS UNIX or UNIX System Services for their normal processing. Plan your configuration if you do not want to use the default UNIX System Services configuration.

No special action or customization is necessary in order for IBM MQ to use UNIX services as long as a system-wide default OMVS segment has been set up.

Users who do not want IBM MQ to invoke UNIX System Services, using the guest or default UID and OMVS segment, need only model a new OMVS segment based on the default segment, as IBM MQ requires no special permissions, and does not run within UNIX as a superuser.

Planning your z/OS TCP/IP environment

To get the best throughput through your network, you must use TCP/IP send and receive buffers with a size of 64 KB, or greater. With this size, the system optimizes its buffer sizes.

See Dynamic right sizing for high latency networks for more information.

You can check your system buffer size by using the following Netstat command, for example:

```
TSO NETSTAT ALL (CLIENT csq1CHIN
```

The results display much information, including the following two values:

```
ReceiveBufferSize: 0000065536  
SendBufferSize: 0000065536
```

65536 is 64 KB. If your buffer sizes are less than 65536, you must work with your network team to increase the **TCPSENDBFRSIZE** and **TCPRCVBUFRSIZE** values in the PROFILE DDName in the TCPIP procedure. For example, you might use the following command:

```
TCPCONFIG TCPSENDBFRSIZE 65536 TCPRCVBUFRSIZE 65536
```

If you are unable to change your system-wide **TCPSENDBFRSIZE** or **TCPRCVBUFRSIZE** settings, contact your IBM Software Support center.

Installing and uninstalling

Before you start installing IBM MQ, consider how you want to use it. Use these topics to help you to prepare for installation, install the product, and verify the installation. There is also information to help you to uninstall the product.

UNIX **Linux** **Windows** Completing the following topics in sequence will help you to correctly install and uninstall IBM MQ and its components on distributed platforms:

1. Planning your installation
2. Checking requirements
3. Preparing your system
4. Installing components
5. Verifying your installation
6. Uninstalling

z/OS For information specific to installing IBM MQ on z/OS systems, see “Installing IBM MQ for z/OS” on page 409.

You can also apply and remove maintenance to IBM MQ. See Maintenance tasks in the Migrating and upgrading section.

Planning your installation

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

Before you start installing, consider how you want to use IBM MQ and review the general planning section. See Planning.

Note that HP Integrity NonStop Server uses a different installer from the other UNIX platforms, and a multi-installation layout.

There are a number of choices you must make before you start installing:

When you have finished planning your installation, check your system requirements. See Checking requirements.

Choosing an installation name

Each installation of IBM MQ on UNIX, Linux, and Windows, has a unique identifier known as an installation name. The installation name is used to associate things such as queue managers and configuration files with an installation.

You can choose the installation name and make it meaningful to you. For example, you might call a test system *testMQ*.

If you do not specify an installation name when the product is installed, a default installation name is automatically assigned. For the first installation, this name is *Installation1*. For the second installation, the name is *Installation2*, and so on. The installation name *Installation0* is reserved for an installation of IBM WebSphere MQ Version 7.0.1. The installation name cannot be changed after the product is installed.

On UNIX and Linux systems, the first IBM MQ installation is automatically given an installation name of *Installation1*. For subsequent installations, you can use the **crtmqinst** command to set the installation name before installing the product.

On Windows systems, you can choose the installation name during the installation process.

The installation name can be up to 16 bytes and must be a combination of alphabetic and numeric characters in the ranges a-z, A-Z, and 0-9. You cannot use blank characters. The installation name must be unique, regardless of whether uppercase or lowercase characters are used. For example, the names INSTALLATIONNAME and InstallationName are not unique.

You can find out what installation name is assigned to an installation in a particular location using the **dspmqinst** command.

Installation descriptions

Each installation can also have an installation description. This description can give more detailed information about an installation in cases where the installation name cannot provide enough information. These descriptions can be up to 64 single-byte characters, or 32 double-byte characters. The default installation description is blank. You can set the installation description using the **setmqinst** command.

Related concepts:

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“Choosing an installation location” on page 192

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the *MQ_INSTALLATION_PATH*.

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

Related information:

dspmqinst

setmqinst

crtmqinst

Multiple installations

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

You can choose where each copy of IBM MQ is installed, but each copy must be in a separate installation location. A maximum of 128 installations of IBM MQ can exist on a single machine at a time. One installation can be an installation of IBM WebSphere MQ Version 7.0.1 at fix pack 6, or later. You have a choice:

- Keep the simplicity of maintaining and managing a single installation of IBM MQ on a machine.
- Take advantage of the flexibility that is offered by enabling multiple IBM MQ installations.

Decisions to make before installing

Before you install multiple copies of IBM MQ, you must make several decisions:

Will you have a copy of IBM WebSphere MQ Version 7.0.1 on the system?

When IBM WebSphere MQ Version 7.0.1 at fix pack 6, or later, is installed on the system, there are a number of restrictions to consider:

- On UNIX and Linux systems, IBM WebSphere MQ Version 7.0.1 must be installed in the default location.
- IBM WebSphere MQ Version 7.0.1 must be the first installation on a system. You cannot install IBM WebSphere MQ Version 7.0.1 after installing Version 7.1, or later. If you uninstall version Version 7.0.1, it cannot be reinstalled while a later version of IBM MQ is installed.
- IBM WebSphere MQ Version 7.0.1 is automatically the primary installation. You cannot select another installation as the primary installation while IBM WebSphere MQ Version 7.0.1 is installed.

Where will you install each copy of IBM MQ ?

You can choose the installation location for your installations at Version 7.1, or later. For more information, see “Choosing an installation location” on page 192.

Do you need a primary installation?

A primary installation is an installation to which system-wide locations refer.

For more information, see “Choosing a primary installation” on page 184.

How will your applications connect?

You need to consider how your applications locate the appropriate IBM MQ libraries. For more information, see Connecting applications in a multiple installation environment, and Connecting .NET applications in a multiple installation environment.

Do your existing exits need changing?

If IBM MQ is not installed in the default location, your exits need to be updated. For more information, see Writing exits and installable services on UNIX, Linux and Windows .

Which queue manager will be associated with which installation?

Each queue manager is associated with a particular installation. The installation that a queue manager is associated with limits that queue manager so that it can be administered only by commands from that installation. For more information, see Associating a queue manager with an installation.

How will you set up your environment to work with each installation?

With multiple installations on a system, you need to consider how you will work with particular installations, and how you will issue commands from that installation. You can either specify the full path to the command, or you can use the **setmqenv** or **crtmqenv** command to set environment variables. Setting the environment variables allows you to omit the path to the commands for that installation. For more information, see **setmqenv**, and **crtmqenv**.

When you have answered these questions, you can install IBM MQ using the steps that are provided in “Installing IBM MQ” on page 249.

If you have existing installations of IBM MQ and you want to use the multiple installation capability to migrate from one version of IBM MQ to another version, see Multi-installation queue manager coexistence on UNIX, Linux, and Windows , unless you are using HP Integrity NonStop Server. In the case of HP Integrity NonStop Server, see “Planning your installation on HP Integrity NonStop Server” on page 220 instead.

The IBM message service client for .NET support pack and multiple installations

For multiple version support, on IBM WebSphere MQ Version 7.1 or later, the *Java and .NET Messaging and Web Services* feature must be installed with the IBM MQ product. For more information about installing the .NET feature, see Installing IBM MQ classes for .NET.

Related information:

Configuring multiple installations

Finding installations of IBM MQ on a system

UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version

UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version

Choosing a primary installation

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Before IBM WebSphere MQ Version 7.1, only one instance of the product could be installed at any time. On Windows systems, several global environment variables were set to point to that installation. On UNIX and Linux systems, symbolic links were added to `/usr/lib`, `/usr/bin`, and `/usr/include`, also pointing at that single installation.

From Version 7.1, you can install multiple versions of IBM MQ on UNIX, Linux, and Windows. It is possible to have more than one installation of IBM MQ on one of these systems at any time and, optionally, to configure one of these installations as the primary installation. Environment variables and symbolic links pointing to a single installation are less meaningful when multiple versions exist. However, some functions require these system-wide locations to work. For example, custom user scripts for administering IBM MQ, and third party products. These functions work only on the primary installation.

On UNIX and Linux systems, if you set an installation as the primary installation, symbolic links to the external libraries and control commands of that installation are added into `/usr/lib`, and `/usr/bin`. If you do not have a primary installation, the symbolic links are not created. For a list of the symbolic links that are made to the primary installation, see “External library and control command links to primary installation on UNIX and Linux” on page 188.

On Windows systems, the global environmental variables point to the directories into which the primary installation was installed. These environment variables are used to locate IBM MQ libraries, control commands, and header files. Additionally, on Windows systems, some features of the operating system require the central registration of interface libraries that are then loaded into a single process. With multiple versions of IBM MQ, there would be conflicting sets of IBM MQ libraries. The features would try to load these conflicting sets of libraries into a single process. Therefore, such features can be used only with the primary installation. For details about some of the features that are limited to use with the primary installation, see “Features that can be used only with the primary installation on Windows” on page 191.

If you have an installation of IBM WebSphere MQ Version 7.0.1 on the system, this installation is automatically the primary installation. The primary installation cannot be changed while Version 7.0.1 is installed. If all the installations on the system are at Version 7.1, or later, you can choose whether to have a primary installation. Consider the options in Table 28.

Table 28. Primary installation options

Options	Valid installation configurations		More information
	Primary	Non-primary	
Single installation of Version 7.1, or later.	Version 7.1, or later.	None	If you want to continue working with a single installation in the same way as previous releases, configure your installation as the primary installation. For information about this option, see Single installation of IBM WebSphere MQ Version 7.1, or later, configured as the primary installation
	None	Version 7.1, or later.	If you want to continue working with a single installation, but do not want symbolic links or global environment variables created for you, configure your installation as non-primary. For information about the implications of this option, see Single installation of IBM WebSphere MQ Version 7.1, or later, configured as non-primary
Multiple installations: Version 7.0.1 and Version 7.1, or later.	Version 7.0.1	Version 7.1, or later.	If you want to have multiple installations of IBM MQ, with one at Version 7.0.1, the Version 7.0.1 installation is automatically the primary installation. While IBM WebSphere MQ Version 7.0.1 is installed, you cannot change which installation is the primary installation. For information about this option and its implications, see Multiple installations of IBM MQ, one at Version 7.0.1
Multiple installations: Version 7.1, or later.	Version 7.1, or later.	Version 7.1, or later.	If you want to have multiple installations of IBM MQ at Version 7.1 or greater, you can choose whether to make one of the installations primary. For information about this option, see Multiple installations of IBM WebSphere MQ Version 7.1, or later
	None	Version 7.1, or later.	

Related concepts:

“Choosing an installation location” on page 192

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the `MQ_INSTALLATION_PATH`.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Choosing an installation name” on page 182

Each installation of IBM MQ on UNIX, Linux, and Windows, has a unique identifier known as an installation name. The installation name is used to associate things such as queue managers and configuration files with an installation.

Related information:

Single installation of IBM WebSphere MQ Version 7.1, or later, configured as the primary installation
Marking an IBM MQ installation as primary adds symbolic links, or global environment variables to the system so that the IBM MQ commands and libraries used by applications are automatically available with minimum system setup required.

Single installation of IBM WebSphere MQ Version 7.1, or later, configured as non-primary

If you install IBM WebSphere MQ Version 7.1, or later, as non-primary you might have to configure a library path for applications to load IBM MQ libraries. On Windows, some product capabilities are available only when IBM MQ is configured as primary.

Multiple installations of IBM WebSphere MQ Version 7.1, or later

You can choose to have one of the IBM WebSphere MQ Version 7.1 installations configured as the

primary installation. Your choice depends on how applications locate libraries.

Multiple installations of IBM MQ, one at Version 7.0.1

IBM WebSphere MQ Version 7.1, or later, can coexist with IBM WebSphere MQ Version 7.0.1 with some limitations.

Changing the primary installation

Single installation of IBM WebSphere MQ Version 7.1, or later, configured as the primary installation

Marking an IBM MQ installation as primary adds symbolic links, or global environment variables to the system so that the IBM MQ commands and libraries used by applications are automatically available with minimum system setup required.

You decide where to install IBM MQ.

Where possible, configure applications and scripts to use the system search path to find the IBM MQ control commands or IBM MQ libraries. This configuration of applications and scripts provides maximum flexibility for undertaking future tasks such as migrating to the next release of IBM MQ, or installing a second installation. For more information about options for connecting your applications, see [Connecting applications in a multiple installation environment](#).

On Windows, the first installation is automatically configured as the primary installation. On UNIX and Linux platforms, the first installation onto a system must be manually configured to be the primary installation. Set the primary installation using the `setmqinst` command. For more information, see [Uninstalling, upgrading, and maintaining the primary installation](#).

Related information:

[Changing the primary installation](#)

[Choosing an installation location](#)

[Planning your installation](#)

[Choosing an installation name](#)

Single installation of IBM WebSphere MQ Version 7.1, or later, configured as non-primary

If you install IBM WebSphere MQ Version 7.1, or later, as non-primary you might have to configure a library path for applications to load IBM MQ libraries. On Windows, some product capabilities are available only when IBM MQ is configured as primary.

UNIX and Linux systems

The implications of running a non-primary installation on UNIX and Linux are:

- Applications that locate their IBM MQ libraries using an embedded library path, for example, RPATH, cannot find those libraries if the following conditions are true:
 - IBM MQ is installed in a different directory from the directory specified in the RPATH
 - There are no symbolic links in /usr
- Where applications locate their libraries using an external library path, for example, LD_LIBRARY_PATH, you must configure the external library path to include the `MQ_INSTALLATION_PATH/lib` or `MQ_INSTALLATION_PATH/lib64` directory. The `setmqenv` and `crtmqenv` commands can configure a number of environment variables in the current shell, including the external library path.
- Most IBM MQ processes run as `setuid/setgid`. As a result, when loading user exits they ignore the external library path. User exits that reference IBM MQ libraries can find those libraries only if they are found in the library path embedded within them. They would be resolved if there were a symbolic link in /usr. User exits that are intended to be run on IBM WebSphere MQ Version 7.1, or later can now be built so that they do not refer to IBM MQ libraries at all. Instead they rely on IBM MQ to pass in

function pointers to the IBM MQ functions that the exit can then use. For more information, see Writing exits and installable services on UNIX, Linux and Windows .

For more information about options for connecting your applications, see Connecting applications in a multiple installation environment.

On UNIX and Linux platforms, the first installation onto a system is not automatically configured as the primary installation. However, a single symbolic link is included in `/usr/bin` to locate the **dspmqver** command. If you do not want any symbolic links, you must remove this link using the following command:

```
setmqinst -x -p MQ_INSTALLATION_PATH
```

Windows systems

The implications of running a non-primary installation on Windows are:

- Applications normally find their libraries using the external library path, `PATH`. There is no concept of an embedded library path or explicit library location. If the installation is non-primary, the global `PATH` environment variable does not contain the IBM MQ installation directory. For applications to find IBM MQ libraries, update the `PATH` environment variable to reference the IBM MQ installation directory. The **setmqenv** and **crtmqenv** commands can configure a number of environment variables in the current shell, including the external library path.
- Some product capabilities are available only when an installation is configured as the primary installation; see Features that can be used only with the primary installation on Windows .

By default, on Windows, the first installation is automatically configured as primary. You must manually deselect it as the primary installation.

Related information:

Changing the primary installation

Choosing an installation location

Planning your installation

setmqenv

crtmqenv

Choosing an installation name

Multiple installations of IBM WebSphere MQ Version 7.1, or later

You can choose to have one of the IBM WebSphere MQ Version 7.1 installations configured as the primary installation. Your choice depends on how applications locate libraries.

The IBM MQ libraries, such as `mqm`, which ship with IBM WebSphere MQ Version 7.1 automatically use libraries of the level required by the queue manager to which they are connecting. This means that provided an application locates its IBM MQ libraries from a IBM WebSphere MQ Version 7.1 installation, it can connect to any queue manager on that system. Having one IBM WebSphere MQ Version 7.1 installation configured as primary ensures that if the application finds its IBM MQ interface library, the application can connect to any queue manager.

For more information about connecting applications in a multiple installation environment, see Connecting applications in a multiple installation environment.

The primary installation is not automatically changed when you uninstall the primary installation. If you want another installation to be the primary installation, you must manually set the primary installation using the **setmqinst** command. For more information, see Uninstalling, upgrading, and maintaining the primary installation.

Related information:

Changing the primary installation

Choosing an installation location

Multiple installations

Planning your installation

Choosing an installation name

Multiple installations of IBM MQ, one at Version 7.0.1

IBM WebSphere MQ Version 7.1, or later, can coexist with IBM WebSphere MQ Version 7.0.1 with some limitations.

- On UNIX and Linux systems, Version 7.0.1 can be installed only in a fixed, default location, so you cannot install Version 7.1, or later, in that default location.
- IBM WebSphere MQ Version 7.0.1 is automatically configured as the primary installation. On UNIX and Linux systems, symbolic links are automatically created to the appropriate IBM MQ directories. On Windows, everything that the product provided is registered globally. IBM WebSphere MQ Version 7.0.1 must be installed in this way to work. So where IBM WebSphere MQ Version 7.0.1 is installed, a IBM WebSphere MQ Version 7.1, or later, installation cannot be made primary.

The libraries from IBM WebSphere MQ Version 7.1, or later, can work with any queue manager running under IBM WebSphere MQ Version 7.0.1, or later. If an application needs to connect to queue managers running under Version 7.0.1 as well as later versions, it can continue to operate normally if the following conditions are true:

- It locates IBM WebSphere MQ Version 7.1, or later, libraries at run time.
- It uses only functions available in Version 7.0.1.

For more information about connecting applications in a multiple installation environment, see [Connecting applications in a multiple installation environment](#).

The primary installation is not automatically changed when you uninstall IBM WebSphere MQ Version 7.0.1. If you want another installation to be the primary installation, you must manually set the primary installation using the **setmqinst** command. For more information, see [Uninstalling, upgrading, and maintaining the primary installation](#).

Related information:

Choosing an installation location

Planning your installation

Multiple installations

Choosing an installation name

External library and control command links to primary installation on UNIX and Linux

On UNIX and Linux platforms the primary installation is the one to which links from the `/usr` file system are made. However, only a subset of those links created with previous releases are now made.

No links are created from `/usr/include` to any installation and only links to external libraries and documented control commands are made from `/usr/lib`, and where appropriate, `/usr/lib64` (external libraries) and `/usr/bin` (control commands).

In order to run these commands you must complete the following steps:

1. provide a full path to the command in an available IBM MQ installation,
2. use the `setmqenv` script to update your shell environment,
3. manually add the `bin` directory from an IBM MQ installation directory to your `PATH`,

4. run the **setmqinst** command as root to make one of your existing IBM MQ installations the primary installation.

External libraries

Links are made to the following external libraries, both 32-bit and 64-bit:

- libmqm
- libmqm_r
- libmqmxa
- libmqmxa_r
- libmqmax
- libmqmax_r
- libmqmcb
- libmqmcb_r
- libmqic
- libmqic_r
- libmqcxa
- libmqcxa_r
- libmqicb
- libmqicb_r
- libimqb23ia
- libimqb23ia_r
- libimqc23ia
- libimqc23ia_r
- libimqs23ia
- libimqs23ia_r
- libmqmzf
- libmqmzf_r

The following 64-bit only libraries are also linked to:

- libmqmxa64
- libmqmxa64_r
- libmqcxa64
- libmqcxa64_r

Control commands

The following control commands are linked to from `/usr/bin`:

- addmqinf
- amqcrs6a
- amqcrsta
- amqmfscck
- crtmqinst
- dltnmqinst
- dspmqinst
- setmqinst
- crtmqcvx

- crtmqm
- dltmqm
- dmpmqaut
- dmpmqlog
- dspmq
- dspmqaut
- dspmqcsv
- dspmqfls
- dspmqinf
- dspmqrte
- dspmqtrc
- dspmqtrn
- dspmqver
- endmqcsv
- endmqlsr
- endmqm
- endmqtrc
- rcdmqimg
- rcrmobj
- rmvmqinf
- rsvmqtrn
- runmqchi
- runmqchl
- runmqckm
- runmqdlq
- runmqlsr
- runmqsc
- runmqtmc
- runmqtrm
- setmqaut
- setmqenv
- setmqm
- setmqprd
- strmqcsv
- strmqikm
- strmqm
- strmqtrc

Related concepts:

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“Features that can be used only with the primary installation on Windows”

Some Windows operating-system features can be used only with the primary installation. This restriction is due to the central registration of interface libraries, which might conflict as a result of multiple versions of IBM MQ being installed.

Features that can be used only with the primary installation on Windows

Some Windows operating-system features can be used only with the primary installation. This restriction is due to the central registration of interface libraries, which might conflict as a result of multiple versions of IBM MQ being installed.

The .NET monitor

The IBM MQ .NET monitor can run in two different modes: transactional and non-transactional. The transactional mode uses MSDTC transaction coordination and requires that the .NET monitor is registered with COM+. The .NET monitor from the primary installation is the only .NET monitor that is registered with COM+.

Any attempt to run the .NET monitor in transactional mode with a non-primary installation results in the failure of the .NET monitor to enlist with MSDTC. The .NET monitor receives an MQRC_INSTALLATION_MISMATCH error, which in turn results in an AMQ8377 error message on the console.

COM/ActiveX interface classes

The COM/ActiveX interface classes are registered only for the primary installation. If there is an installation of IBM WebSphere MQ Version 7.0.1 on the system, the COM/ActiveX interface classes registered are not capable of connecting to queue managers running under other installations. If the primary installation is an installation of IBM WebSphere MQ Version 7.1 or later, the interface classes can connect to queue managers associated with any installation. Server COM/ActiveX applications are limited by this restriction, but client applications can connect to any queue manager.

Any attempt to start a COM/ActiveX application that uses libraries from installations other than the primary installation results in failure with an MQRC_Q_MGR_NOT_AVAILABLE error.

Related concepts:

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“External library and control command links to primary installation on UNIX and Linux” on page 188

On UNIX and Linux platforms the primary installation is the one to which links from the /usr file system are made. However, only a subset of those links created with previous releases are now made.

Uninstalling, upgrading, and maintaining the primary installation

On all platforms, if you uninstall the primary installation, it ceases to be the primary installation. You must run the **setmqinst** command to select a new primary installation. On Windows, if you update the primary installation, it continues to be the primary installation. If you apply a fix pack to the primary installation, it continues to be the primary installation.

Be cautious about the effect uninstalling or upgrading the primary installation has on applications. Applications might be using the linkage library of the primary installation to switch to the linkage library

of another installation. If such an application is running, you might not be able to uninstall the primary installation. The operating system might have locked the link library of the primary installation on behalf of the application. If the primary installation has been uninstalled, an application that loads the IBM MQ libraries it requires by linking to the primary installation is not able to start.

The solution is to switch the primary installation to another installation before uninstalling. Stop, and restart applications that are linked through the previous primary installation before uninstalling it.

Windows

If you update the primary installation, it stops being the primary installation at the beginning of the update procedure. If, by the end of the update procedure, you have not made another installation primary, the upgraded installation is made primary again.

Maintenance

If you apply a fix pack to the primary installation, it stops being the primary installation at the beginning of the maintenance procedure. If, by the end of the maintenance procedure, you have not made another installation primary, the upgraded installation is made primary again.

Related concepts:

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“Uninstalling” on page 389

The topics in this section provide instructions on how to uninstall components.

Related information:

Changing the primary installation

IBM MQ maintenance tasks








Choosing an installation location

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the `MQ_INSTALLATION_PATH`.

Default location

The default location for the IBM MQ product code is shown in the following table:

Table 29. Installation location of IBM MQ

Platform	Installation location
 AIX	/usr/mqm
   Linux , HP-UX, and Solaris	/opt/mqm
 IBM i	/QIBM/ProdData/mqm
 Windows systems	C:\Program Files\IBM\WebSphere MQ
 Windows data directories	C:\ProgramData\IBM\MQ

Important: For Windows installations, the directories are as stated, unless there is a previous installation of IBM MQ that still contains registry entries or queue managers, or both. In this situation, the new installation uses the old data directory location. For more information, see Program and data directory locations.

A Windows 32-bit client installs on a 32-bit machine, by default, only into C:\Program Files\IBM\WebSphere MQ.

A Windows 64-bit client installs on a 64-bit machine, by default, only into C:\Program Files\IBM\WebSphere MQ.

On IBM i, IBM MQ can only be installed in the default location. For more information about the directory structure of IBM i, see Directory structure on IBM i

On UNIX and Linux systems, working data is stored in /var/mqm, but you cannot change this location. For more information about the directory structure of UNIX and Linux systems, see Directory structure on UNIX and Linux systems.

Custom location installation

For an installation into a custom location, the path specified must either be an empty directory, or a path that does not exist. The length of the path is limited to 256 bytes.

- On UNIX and Linux systems, the path must not contain spaces.
- On AIX, the product is installed into a User Specified Install Location (USIL), which can be either an existing USIL or a new USIL that is automatically created by the installation process. If a custom location is specified, the product location is the path specified during installation, plus /usr/mqm.

For example, the path specified is /usr/custom_location. The *MQ_INSTALLATION_PATH* is /usr/custom_location/usr/mqm.

Access permissions for the USIL directory should be set to rwx for user and r-x for group and others (755).

- On Windows, Linux, HP-UX, and Solaris, the product location is the same path as specified during installation.
For example, on Linux, the path specified is /opt/custom_location. The *MQ_INSTALLATION_PATH* is /opt/custom_location.
- On Linux, Solaris, and HP-UX, you can install IBM MQ into a non empty *MQ_INSTALLATION_PATH* directory.

On Linux and Solaris you do this by setting the environment variable *AMQ_OVERRIDE_EMPTY_INSTALL_PATH* to 1 before starting the installation.

On HP-UX you need to create the file /tmp/AMQ_OVERRIDE_EMPTY_INSTALL_PATH before starting the installation.

Note, that a non empty directory in this context, indicates a directory which contains system files and directories.

For each installation, all of the IBM MQ components that you require must be installed in the same location.

For more information about how to install to a custom location, see the installation topics for the appropriate platform.

Additional location restrictions

New IBM MQ installations should not be located in the following paths:

- In a path that is a subdirectory of another existing installation.

- In a path that is part of the direct path to an existing installation.
- In a path that is a subdirectory of the default location, for example:
 - /usr/mqm on AIX
 - /opt/mqm on Linux, Solaris and HP-UX platforms
- In a directory or subdirectory that is, or might later be used by another product, for example, an IBM Db2 installation, or operating system component.

An installation should not be located in /opt/mqm/v80, /opt/mqm/v75, /opt/mqm/inst2/mq71, or other directory located under /opt/mqm on Linux, Solaris and HP-UX platforms.

If IBM MQ is installed in /opt/IBM/MQ/installations/1, you can not install in /opt/IBM/MQ/installations/1/a. Additionally, you should not install a new installation to /opt/IBM/MQ. However, you can install a new installation in /opt/IBM/MQ/installations/2 or /opt/IBM/MQnew because neither of these is a part of the direct path /opt/IBM/MQ/installations/1.

You must not install to any directory located under /opt/IBM/db2.

The reason an installation should not be located in a path that is a subdirectory of the default location is to avoid the risk if you later decide to install IBM MQ into the default location, and cannot then do so. If you do subsequently install into the default location, because IBM MQ has full access rights over the installation directory, existing files might be replaced or deleted. Scripts that you might subsequently run to uninstall IBM MQ might remove the installation directory at the end of the script.

Related concepts:

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Choosing an installation name” on page 182

Each installation of IBM MQ on UNIX, Linux, and Windows, has a unique identifier known as an installation name. The installation name is used to associate things such as queue managers and configuration files with an installation.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“Choosing what to install”

You can select the components or features that you require when you install IBM MQ.

Related tasks:

“Installing an IBM MQ server” on page 254

After preparing your system for installation you can install IBM MQ by following the appropriate instructions for your platform. After installation, you might want to verify your installation to check that installation has been successful.

Choosing what to install

You can select the components or features that you require when you install IBM MQ.

Important: Ensure that your enterprise has the correct license, or licenses, for the components that you are going to install. See IBM MQ license information for more details.

IBM MQ can be installed as a server or a client. The installation images can be downloaded, or IBM MQ can be installed from a DVD.

An IBM MQ server is an installation of one or more queue managers that provide queuing services to one or more clients. All the IBM MQ objects, for example queues, exist only on the queue manager

machine (the IBM MQ server machine), and not the client. An IBM MQ server can also support local IBM MQ applications. To install an IBM MQ server, see “Installing an IBM MQ server” on page 254.

An IBM MQ MQI client is a component that allows an application running on one system to communicate with a queue manager running on another system. The output from the call is sent back to the client, which passes it back to the application. To install an IBM MQ MQI client, see Installing an IBM MQ client.

It is possible to have both a server and a client installation on the same system. See “Installing an IBM MQ client” on page 319.

IBM MQ Advanced Message Security is a separately installed and licensed component of IBM MQ and is another option on the IBM MQ installer. To install IBM MQ Advanced Message Security, see “Installing IBM MQ Advanced Message Security” on page 353.

For detailed explanations of all the components that you can install, see the following platform-specific topics.

Related concepts:

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Choosing an installation location” on page 192

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the `MQ_INSTALLATION_PATH`.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

“Choosing an installation name” on page 182

Each installation of IBM MQ on UNIX, Linux, and Windows, has a unique identifier known as an installation name. The installation name is used to associate things such as queue managers and configuration files with an installation.

IBM MQ components for AIX systems

You can select the components that you require when you install IBM MQ.

On AIX each component of IBM MQ is represented by a fileset. Table 30 shows the filesets that are available when installing an IBM MQ server or client on an AIX system:

Table 30. IBM MQ filesets for AIX systems

Component	Description	Server media	Client media	Hypervisor	Fileset name
Runtime	Contains files that are common to both server and client installations. Note: This component must be installed.	✓	✓	✓	mqm.base.runtime
Server	You can use the server to run queue managers on your system and connect to other systems over a network. Provides messaging and queuing services to applications, and support for IBM MQ client connections.	✓		✓	mqm.server.rte

Table 30. IBM MQ filesets for AIX systems (continued)

Component	Description	Server media	Client media	Hypervisor	Fileset name
Standard Client	The IBM MQ MQI client is a small subset of IBM MQ, without a queue manager, that uses the queue manager and queues on other (server) systems. It can be used only when the system it is on is connected to another system that is running a full server version of IBM MQ. The client and the server can be on the same system if required.	✓	✓	✓	mqm.client.rte
SDK	The SDK is required for compiling applications. It includes sample source files, and the bindings (files .H, .LIB, .DLL, and others), that you need to develop applications to run on IBM MQ.	✓	✓	✓	mqm.base.sdk
Sample programs	The sample application programs are needed if you want to check your IBM MQ installation using the verification procedures.	✓	✓	✓	mqm.base.samples
Java messaging	The files needed for messaging using Java (includes Java Messaging Service).	✓	✓	✓	mqm.java.rte
Man pages	UNIX man pages, in U.S. English, for: control commands MQI calls MQSC commands	✓	✓	✓	mqm.man.en_US.data
Java JRE	A Java Runtime Environment that is used by those parts of IBM MQ that are written in Java.	✓	✓	✓	mqm.jre.rte
Message Catalogs	For available languages, see the table of message catalogs that follows.	✓	✓	✓	
IBM Global Security Kit	IBM Global Security Kit V8 Certificate and SSL, or TLS, Base Runtime. You must install the Java JRE component to install this component.	✓	✓	✓	mqm.gskit.rte
Telemetry Service	MQ Telemetry supports the connection of Internet Of Things (IOT) devices (that is, remote sensors, actuators and telemetry devices) that use the IBM MQ Telemetry Transport (MQTT) protocol. The telemetry service, which is also know as the MQ Extended Reach (MQXR) service, enables a queue manager to act as an MQTT server, and communicate with MQTT client apps. A set of MQTT client libraries is also available in the free download IBM Messaging Telemetry Clients SupportPac. These libraries help you write the MQTT client apps that IOT devices use to communicate with MQTT servers. See also “Installing IBM MQ Telemetry” on page 214.	✓		✓	mqm.xr.service

Table 30. IBM MQ filesets for AIX systems (continued)

Component	Description	Server media	Client media	Hypervisor	Fileset name
Managed File Transfer	MQ Managed File Transfer transfers files between systems in a managed and auditable way, regardless of file size or the operating systems used. For information about the function of each component, see IBM MQ Managed File Transfer product options.	✓		✓	mqm.ft.agent mqm.ft.base mqm.ft.logger mqm.ft.service mqm.ft.tools
Advanced Message Security	Provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications. You must install this component on all IBM MQ installations that host queues you want to protect. You must install the IBM Global Security Kit component on any IBM MQ installation that is used by a program that puts or gets messages to or from a protected queue, unless you are using only Java client connections. You must install the Java JRE component to install this component.	✓		✓	mqm.ams.rte
V8.0.0.4 AMQP Service	Install this component to make AMQP channels available. AMQP channels support MQ Light APIs. You can use AMQP channels to give AMQP applications access to the enterprise-level messaging facilities provided by IBM MQ.	✓			mqm.amqp.rte

Table 31. IBM MQ message catalogs for AIX systems

Message catalog language	Component name
Brazilian Portuguese	mqm.msg.pt_BR
Czech	mqm.msg.cs_CZ
French	mqm.msg.fr_FR
German	mqm.msg.de_DE
Hungarian	mqm.msg.hu_HU
Italian	mqm.msg.it_IT
Japanese	mqm.msg.ja_JP, mqm.msg.Ja_JP
Korean	mqm.msg.ko_KR
Polish	mqm.msg.pl_PL
Russian	mqm.msg.ru_RU
Spanish	mqm.msg.es_ES
Simplified Chinese	mqm.msg.zh_CN, mqm.msg.Zh_CN
Traditional Chinese	mqm.msg.zh_TW, mqm.msg.Zh_TW
U.S. English	mqm.msg.en_US

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

IBM MQ client components for HP Integrity NonStop Server

There are no optional components within the client installer when you install the IBM MQ client for HP Integrity NonStop Server.

An installation of the IBM MQ client for HP Integrity NonStop Server contains product binary files, command utilities, and samples.

There are no optional components for HP Integrity NonStop Server.

Related concepts:

“Planning your installation on HP Integrity NonStop Server” on page 220

This section describes what to do to prepare your system for installing IBM MQ client for HP Integrity NonStop Server .

“Hardware and software requirements on HP Integrity NonStop Server systems” on page 231

Check that the server environment meets the prerequisites for installing the IBM MQ client for HP Integrity NonStop Server. Check the product readme files and install missing prerequisite software supplied on the server CD.

Related information:

Planning your IBM MQ client environment on HP Integrity NonStop Server

IBM MQ components for HP-UX systems

You can select the components that you require when you install IBM MQ.

Table 32 shows the components that are available when installing an IBM MQ server or client on an HP-UX system:

Table 32. IBM MQ components for HP-UX systems.

Component	Description	Server media	Client media	Component name
Runtime	Contains files that are common to both server and client installations. Note: This component must be installed.	✓	✓	MQSERIES.MQM-RUNTIME
Server	You can use the server to run queue managers on your system and connect to other systems over a network. Provides messaging and queuing services to applications, and support for IBM MQ client connections.	✓		MQSERIES.MQM-SERVER
Standard Client	The IBM MQ MQI client is a small subset of IBM MQ, without a queue manager, that uses the queue manager and queues on other (server) systems. It can be used only when the system it is on is connected to another system that is running a full server version of IBM MQ. The client and the server can be on the same system if required.	✓	✓	MQSERIES.MQM-CL-HPUX

Table 32. IBM MQ components for HP-UX systems (continued).


Component	Description	Server media	Client media	Component name
SDK	The SDK is required for compiling applications. It includes sample source files, and the bindings (files .H, .LIB, .DLL, and others), that you need to develop applications to run on IBM MQ.	✓	✓	MQSERIES.MQM-BASE
Sample programs	The sample application programs are needed if you want to check your IBM MQ installation using the verification procedures.	✓	✓	MQSERIES.MQM-SAMPLES
Java messaging	The files needed for messaging using Java (includes Java Messaging Service).	✓	✓	MQSERIES.MQM-JAVA
Man pages	UNIX man pages, in U.S. English, for: control commands MQI calls MQSC commands	✓	✓	MQSERIES.MQM-MAN
Java JRE	A Java Runtime Environment that is used by those parts of IBM MQ that are written in Java.	✓	✓	MQSERIES.MQM-JAVAJRE
Message Catalogs	For available languages, see the table of message catalogs that follows.	✓	✓	
IBM Global Security Kit	IBM Global Security Kit V8 Certificate and SSL, or TLS, Base Runtime. You must install the Java JRE component to install this component.	✓	✓	MQSERIES.MQM-GSKIT
Managed File Transfer	MQ Managed File Transfer transfers files between systems in a managed and auditable way, regardless of file size or the operating systems used. For information about the function of each component, see IBM MQ Managed File Transfer product options.	✓		MQSERIES.MQM-FTAGENT MQSERIES.MQM-FTBASE MQSERIES.MQM-FTLOGGER MQSERIES.MQM-FTSERVICE MQSERIES.MQM-FTTOOLS
Advanced Message Security	Provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications. You must install this component on all IBM MQ installations that host queues you want to protect. You must install the IBM Global Security Kit component on any IBM MQ installation that is used by a program that puts or gets messages to or from a protected queue, unless you are using only Java client connections. You must install the Java JRE component to install this component.	✓		MQSERIES.MQM-AMS
 V8.0.0.4 AMQP Service	Install this component to make AMQP channels available. AMQP channels support MQ Light APIs. You can use AMQP channels to give AMQP applications access to the enterprise-level messaging facilities provided by IBM MQ.	✓		MQSERIES.MQM-AMQP

Table 33. IBM MQ message catalogs for HP-UX systems.

A two-column table listing the available message catalogs.

Message catalog language	Component name
Brazilian Portuguese	MQSERIES.MQM-MC-PORT
Czech	MQSERIES.MQM-MC-CZECH
French	MQSERIES.MQM-MC-FRENCH
German	MQSERIES.MQM-MC-GERMAN
Hungarian	MQSERIES.MQM-MC-HUNGARIAN
Italian	MQSERIES.MQM-MC-ITALIAN
Japanese	MQSERIES.MQM-MC-JAPAN
Korean	MQSERIES.MQM-MC-KOREAN
Polish	MQSERIES.MQM-MC-POLISH
Russian	MQSERIES.MQM-MC-RUSSIAN
Spanish	MQSERIES.MQM-MC-SPANISH
Simplified Chinese	MQSERIES.MQM-MC-CHINES
Traditional Chinese	MQSERIES.MQM-MC-CHINET
U.S. English	not applicable

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

IBM MQ components for Linux systems

You can select the components that you require when you install IBM MQ.

Table 34 shows the components that are available when installing an IBM MQ server or client on a Linux system:

Table 34. IBM MQ components for Linux systems

Component	Description	Server media	Client media	Hypervisor	Component name
Runtime	Contains files that are common to both server and client installations. Note: MQSeriesRuntime component must be installed.	✓	✓	✓	MQSeriesRuntime
Server	You can use the server to run queue managers on your system and connect to other systems over a network. Provides messaging and queuing services to applications, and support for IBM MQ client connections.	✓		✓	MQSeriesServer

Table 34. IBM MQ components for Linux systems (continued)

Component	Description	Server media	Client media	Hypervisor	Component name
Standard Client	The IBM MQ MQI client is a small subset of IBM MQ, without a queue manager, that uses the queue manager and queues on other (server) systems. It can be used only when the system it is on is connected to another system that is running a full server version of IBM MQ. The client and the server can be on the same system if required.	✓	✓	✓	MQSeriesClient
SDK	The SDK is required for compiling applications. It includes sample source files, and the bindings (files .H, .LIB, .DLL, and others), that you need to develop applications to run on IBM MQ.	✓	✓	✓	MQSeriesSDK
Sample programs	The sample application programs are needed if you want to check your IBM MQ installation using the verification procedures.	✓	✓	✓	MQSeriesSamples
Java messaging	The files needed for messaging using Java (includes Java Messaging Service).	✓	✓	✓	MQSeriesJava
Man pages	UNIX man pages, in U.S. English, for: control commands MQI calls MQSC commands	✓	✓	✓	MQSeriesMan
Java JRE	A Java Runtime Environment that is used by those parts of IBM MQ that are written in Java.	✓	✓	✓	MQSeriesJRE
Message Catalogs	For available languages, see the table of message catalogs that follows.	✓	✓	✓	
IBM Global Security Kit	IBM Global Security Kit V8 Certificate and SSL, or TLS, Base Runtime. You must install the Java JRE component to install this component.	✓	✓	✓	MQSeriesGSKit

Table 34. IBM MQ components for Linux systems (continued)

Component	Description	Server media	Client media	Hypervisor	Component name
Telemetry Service	<p>MQ Telemetry supports the connection of Internet Of Things (IOT) devices (that is, remote sensors, actuators and telemetry devices) that use the IBM MQ Telemetry Transport (MQTT) protocol. The telemetry service, which is also known as the MQ Extended Reach (MQXR) service, enables a queue manager to act as an MQTT server, and communicate with MQTT client apps.</p> <p>The telemetry service is only available on Linux for System x (64 bit) and Linux for IBM Z .</p> <p>A set of MQTT client libraries is also available in the free download IBM Messaging Telemetry Clients SupportPac. These libraries help you write the MQTT client apps that IOT devices use to communicate with MQTT servers.</p> <p>See also “Installing IBM MQ Telemetry” on page 214.</p>	✓		✓	MQSeriesXRService
MQ Explorer	Use MQ Explorer to administer and monitor resources on Linux x86-64 systems. Also available on x86 via a standalone installer from MSOT.	✓		✓	MQSeriesExplorer
Managed File Transfer	MQ Managed File Transfer transfers files between systems in a managed and auditable way, regardless of file size or the operating systems used. For information about the function of each component, see IBM MQ Managed File Transfer product options.	✓		✓	MQSeriesFTAgent MQSeriesFTBase MQSeriesFTLogger MQSeriesFTService MQSeriesFTTools
Advanced Message Security	<p>Provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications. You must install this component on all IBM MQ installations that host queues you want to protect.</p> <p>You must install the IBM Global Security Kit component on any IBM MQ installation that is used by a program that puts or gets messages to or from a protected queue, unless you are using only Java client connections.</p> <p>You must install the Java JRE component to install this component.</p>	✓		✓	MQSeriesAMS

Table 34. IBM MQ components for Linux systems (continued)

Component	Description	Server media	Client media	Hypervisor	Component name
V8.0.0.4 AMQP Service	Install this component to make AMQP channels available. AMQP channels support MQ Light APIs. You can use AMQP channels to give AMQP applications access to the enterprise-level messaging facilities provided by IBM MQ.	✓			MQSeries® AMQP

Table 35. IBM MQ message catalogs for Linux systems

Message catalog language	Component name
Brazilian Portuguese	MQSeriesMsg_pt
Czech	MQSeriesMsg_cs
French	MQSeriesMsg_fr
German	MQSeriesMsg_de
Hungarian	MQSeriesMsg_hu
Italian	MQSeriesMsg_it
Japanese	MQSeriesMsg_ja
Korean	MQSeriesMsg_ko
Polish	MQSeriesMsg_pl
Russian	MQSeriesMsg_ru
Spanish	MQSeriesMsg_es
Simplified Chinese	MQSeriesMsg_Zh_CN
Traditional Chinese	MQSeriesMsg_Zh_TW
U.S. English	not applicable

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

IBM MQ components for Solaris systems

You can select the components that you require when you install IBM MQ.

Table 36 on page 204 shows the components that are available when installing an IBM MQ server or client on a Solaris system.

Note: When you install interactively on Solaris systems, the options that are available install various combinations of the components listed in this table. Details are given in the “Interactive installation” on page 205 section.

Table 36. IBM MQ components for Solaris systems

Component	Description	Server media	Client media	Hypervisor	Component name
Runtime	Contains files that are common to both server and client installations. Note: This component must be installed.	✓	✓		runtime
Server	You can use the server to run queue managers on your system and connect to other systems over a network. Provides messaging and queuing services to applications, and support for IBM MQ client connections.	✓			server
Standard Client	The IBM MQ MQI client is a small subset of IBM MQ, without a queue manager, that uses the queue manager and queues on other (server) systems. It can be used only when the system it is on is connected to another system that is running a full server version of IBM MQ. The client and the server can be on the same system if required.	✓	✓		sol_client
SDK	The SDK is required for compiling applications. It includes sample source files, and the bindings (files .H, .LIB, .DLL, and others), that you need to develop applications to run on IBM MQ.	✓	✓		base
Sample programs	The sample application programs are needed if you want to check your IBM MQ installation using the verification procedures.	✓	✓		samples
Java messaging	The files needed for messaging using Java (includes Java Messaging Service).	✓	✓		java
Man pages	UNIX man pages, in U.S. English, for: control commands MQI calls MQSC commands	✓	✓		man
Java JRE	A Java Runtime Environment that is used by those parts of IBM MQ that are written in Java.	✓	✓		jre
Message Catalogs	For available languages, see the table of message catalogs that follows.	✓	✓		
IBM Global Security Kit	IBM Global Security Kit V8 Certificate and SSL, or TLS, Base Runtime. You must install the Java JRE component to install this component.	✓	✓		gskit
Managed File Transfer	MQ Managed File Transfer transfers files between systems in a managed and auditable way, regardless of file size or the operating systems used. For information about the function of each component, see IBM MQ Managed File Transfer product options.	✓			ftagent ftbase ftlogger ftservice fttools

Table 36. IBM MQ components for Solaris systems (continued)

Component	Description	Server media	Client media	Hypervisor	Component name
Advanced Message Security	Provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications. You must install this component on all IBM MQ installations that host queues you want to protect. You must install the IBM Global Security Kit component on any IBM MQ installation that is used by a program that puts or gets messages to or from a protected queue, unless you are using only Java client connections. You must install the Java JRE component to install this component.	✓			mqams
V8.0.0.4 AMQP Service	Install this component to make AMQP channels available. AMQP channels support MQ Light APIs. You can use AMQP channels to give AMQP applications access to the enterprise-level messaging facilities provided by IBM MQ.	✓			amqp

Table 37. IBM MQ message catalogs for Solaris systems.

A two-column table listing the available message catalogs.

Message catalog language	Component name
Brazilian Portuguese	Pt_BR
Czech	Cs_CZ
French	Fr_FR
German	De_DE
Hungarian	Hu_HU
Italian	It_IT
Japanese	Ja_JP
Korean	Ko_KR
Polish	Pl_PL
Russian	Ru_RU
Spanish	Es_ES
Simplified Chinese	Zh_CN
Traditional Chinese	Zh_TW
U.S. English	not applicable

Interactive installation

The options available with interactive installation install various combinations of the product components described in the previous tables. The following table shows you what will be installed for each option, together with the option number on the server and client DVDs:

Table 38. IBM MQ interactive installation options for Solaris systems.

A four-column table listing interactive installation options and the components installed with each one. Server and client option numbers are also listed.

Interactive installation option	Components installed	Server DVD option number	Client DVD option number
IBM MQ Server	base runtime server java gskit	1	
Man pages	runtime man	2	1
Sample programs	base runtime samples	3	2
IBM MQ MQI client libraries (including Java, JMS, and Web Services support)	base runtime sol_client java gskit	4	3
IBM Java runtime for Solaris, Java 2 Technology Edition, Version 6	jre runtime	5	
IBM Global Security Kit for IBM MQ	gskit jre runtime	6	
IBM MQ Managed File Transfer Service	ftservice ftbase jre java runtime ftagent	7	
IBM MQ Managed File Transfer Tools	fttools ftbase jre java runtime	8	
IBM MQ Managed File Transfer Agent	ftagent ftbase jre java runtime	9	
IBM MQ Managed File Transfer Logger	ftlogger ftbase jre java runtime server	10	
IBM MQ Advanced Message Security	runtime mqams	11	
Spanish message catalog	runtime Es_ES	12	4

Table 38. IBM MQ interactive installation options for Solaris systems (continued).

A four-column table listing interactive installation options and the components installed with each one. Server and client option numbers are also listed.

Interactive installation option	Components installed	Server DVD option number	Client DVD option number
French message catalog	runtime Fr_FR	13	5
German message catalog	runtime De_DE	14	6
Japanese message catalog	runtime Ja_JP	15	7
Italian message catalog	runtime It_IT	16	8
Brazilian Portuguese message catalog	runtime Pt_BR	17	9
Traditional Chinese message catalog	runtime Zh_TW	18	10
Simplified Chinese message catalog	runtime Zh_CN	19	11
Korean message catalog	runtime Ko_KR	20	12
Russian message catalog	runtime Ru_RU	21	13
Hungarian message catalog	runtime Hu_HU	22	14
Polish message catalog	runtime Pl_PL	23	15
Czech message catalog	runtime Cs_CZ	24	16

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

IBM MQ features for Windows systems

You can select the features that you require when you install IBM MQ.


The following table shows the features that are available when installing an IBM MQ server or client on a Windows system.

Interactive displayed name	Non-interactive displayed name	Description	Server media	Client media
Server	Server	You can use the server to run queue managers on your system and connect to other systems over a network. Provides messaging and queuing services to applications, and support for IBM MQ client connections.	✓	
MQ Explorer	Explorer	MQ Explorer allows you to administer and monitor resources in IBM MQ.	✓	
Managed File Transfer Service	MFT_Service	<p>The IBM MQ Managed File Transfer Service install option installs a file transfer agent that has additional capabilities beyond those provided by the file transfer agent installed via the IBM MQ Managed File Transfer Agent install option. These additional capabilities are:-</p> <ul style="list-style-type: none"> • Create protocol bridge agents which are used to send and receive files with legacy FTP, FTPS or SFTP servers • Deploy the Web Gateway feature which provides RESTful interfaces for building web applications that transfer files <p>The IBM MQ Managed File Transfer Service install option must be installed on systems where the IBM MQ Server install option is already installed.</p>	✓	

Interactive displayed name	Non-interactive displayed name	Description	Server media	Client media
Managed File Transfer Logger	MFT_Logger	The IBM MQ Managed File Transfer Logger install option installs a file transfer logger which connects to an IBM MQ queue manager, often the queue manager designated as the coordination queue manager. It logs file transfer audit related data to either a database or a file. It must be installed on systems where the IBM MQ Server install option is already installed.	✓	
Managed File Transfer Agent	MFT_Agent	The IBM MQ Managed File Transfer Agent install option installs a file transfer agent which connects to an IBM MQ queue manager and transfers file data, as messages, to other file transfer agents. These must be installed either as part of the IBM MQ Managed File Transfer Agent or IBM MQ Managed File Transfer Service install options.	✓	

Interactive displayed name	Non-interactive displayed name	Description	Server media	Client media
Managed File Transfer Tools	MFT_Tools	The IBM MQ Managed File Transfer Tools install option installs command line tools that are used to interact with file transfer agents. You can use these tools to start file transfers, schedule file transfers and create resource monitors from the command line. The IBM MQ Managed File Transfer Tools can be installed and used on either a system where file transfer agents are installed, or on a system where no file transfer agents are installed.	✓	
Local clients\Windows NT Client	Client	The IBM MQ client is a small subset of IBM MQ, without a queue manager, that uses the queue manager and queues on other (server) systems. It can be used only when the system it is on is connected to another system that is running a full server version of IBM MQ. The client and server can be on the same system if required.	✓	✓
Java and .NET Messaging and Web Services	JavaMsg	The files needed for messaging using Java (includes Java Message Service support) and IBM MQ Web Services.	✓	✓

Interactive displayed name	Non-interactive displayed name	Description	Server media	Client media
Development Toolkit	Toolkit	<p>This feature includes sample source files, and the bindings (files .H, .LIB, .DLL, and others), that you need to develop applications to run on IBM MQ. Bindings and samples are provided for the following languages: C, C++, Visual Basic, ActiveX, Cobol, and .NET (including C#). Java and Java Message Service support is included and samples are provided for MTS (COM+), and MQSC.</p>	✓	✓
Telemetry Service	XR_Service	<p>MQ Telemetry supports the connection of Internet Of Things (IOT) devices (that is, remote sensors, actuators and telemetry devices) that use the IBM MQ Telemetry Transport (MQTT) protocol. The telemetry service, which is also know as the MQ Extended Reach (MQXR) service, enables a queue manager to act as an MQTT server, and communicate with MQTT client apps.</p> <p>A set of MQTT client libraries is also available in the free download IBM Messaging Telemetry Clients SupportPac. These libraries help you write the MQTT client apps that IOT devices use to communicate with MQTT servers.</p> <p>See also “Installing IBM MQ Telemetry” on page 214.</p>	✓	

Interactive displayed name	Non-interactive displayed name	Description	Server media	Client media
Advanced Message Security	AMS	<p>Provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications. You must install this component on all IBM MQ installations that host queues you want to protect.</p> <p>You must install the IBM Global Security Kit component on any IBM MQ installation that is used by a program that puts or gets messages to or from a protected queue, unless you are using only Java client connections.</p> <p>You must install the Java JRE component to install this component.</p>	✓	
	AMQP_Service	<p>Install this component to make AMQP channels available. AMQP channels support MQ Light APIs. You can use AMQP channels to give AMQP applications access to the enterprise-level messaging facilities provided by IBM MQ.</p>	✓	

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

IBM MQ components for IBM i

The components are as follows:

Server (Base)

Support to enable you to create and support your own applications. This includes the runtime component that provides support for external applications. It also includes support for client connections from IBM MQ installations on other computers.

Command Reference

Help for the CL commands is provided in HTML format and installed with the product in the /QIBM/ProdData/mqm/doc directory.

Samples (Option 1)

Sample application programs. The source is supplied in the QMQMSAMP library and executable files are supplied in the QMQM library.

AMS (Option 2)

The AMS component.

Documentation

The full product documentation is supplied on the IBM MQ Documentation CD.

Readme file

Latest information about the product that became available after publication of this product documentation or the full documentation. You can find the readme file in the root of the product or documentation CD. Review it before starting to install IBM MQ for IBM i.

IBM MQ Managed File Transfer (MFT) components***BASE**

Support to enable you to create and support your own MFT applications. It also includes support for client connections from IBM MQ MFT installations on other computers.

2 Tools support

3 Agent

4 Services

You must install **BASE* first because the other three options depend on **BASE*. Note that option 4 requires that option 3 is installed.

Related concepts:

“Choosing what to install” on page 194

You can select the components or features that you require when you install IBM MQ.

“IBM MQ features for Windows systems” on page 207

You can select the features that you require when you install IBM MQ.

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

Installing IBM MQ Telemetry

From IBM WebSphere MQ Version 7.1, IBM MQ Telemetry is a component of the main IBM MQ product, and is no longer a separate plug-in. You can choose to install IBM MQ Telemetry when you first install IBM MQ, or when you modify an existing IBM MQ installation. To get a set of client libraries that help you write messaging applications for telemetry, download the free IBM Messaging Telemetry Clients SupportPac.

If IBM WebSphere MQ Version 7.0.1 is installed with the IBM MQ Telemetry plug-in, you must migrate the installation to the latest version of the product. See *Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0* for information and instructions about this process.

IBM MQ Telemetry runtime

The IBM MQ Telemetry runtime enables messaging applications to connect to queue managers across the IBM MQ Telemetry Transport (MQTT) protocol. The MQTT protocol is mainly used for connecting to telemetry and other Internet of Things devices, reaching out to the edge of the network and across fragile communications links.

The core component of the IBM MQ Telemetry runtime is the IBM MQ Extended Reach (MQXR) telemetry service. This service extends a queue manager and enables it to communicate across the IBM MQ Telemetry Transport (MQTT) protocol. To run the telemetry (MQXR) service, you need the following software components:

- A Java Runtime Environment (JRE), for each system on which you run applications
- IBM MQ Java

Support for MQ Explorer

You can use MQ Explorer to configure and manage the IBM MQ Telemetry runtime component. For a queue manager to accept connections from a telemetry device, one or more telemetry channels are needed. To enable MQTT, there is a `define sample` configuration wizard that can be run from MQ Explorer. The wizard runs through a series of steps including defining and starting the telemetry (MQXR) service, setting up the default transmission queue, and configuring a telemetry channel. For more information about using the `define sample` configuration wizard, and any implications, see “*Verifying the installation of IBM MQ Telemetry by using MQ Explorer*” on page 385.

The MQ Explorer support provides the following capabilities:

- Telemetry node and content panel - providing welcome information, `define sample` configuration wizard, run MQTT client utility, Help on IBM MQ Telemetry, and status information about the IBM MQ Telemetry Service.
- Define sample configuration wizard - quickly configures a queue manager to support MQTT.
- New Telemetry Channel wizard - gathers information required to create a telemetry channel object.
- Telemetry Channels node and content panel - displays telemetry channels in the MQ Explorer Content view.

- Telemetry Channel Status node and content panel - displays telemetry channel status in the MQ Explorer Content view.
- MQTT Client Utility - provides a simple GUI for publishing and subscribing to topics.
- Help on IBM MQ Telemetry.

You can install the IBM MQ Telemetry runtime component on one system and configure and manage it using the MQ Explorer installed on another system. However, the components can be installed only on systems with the appropriate prerequisites. For information about these prerequisites, see IBM MQ system requirements.

You can administer IBM MQ Telemetry Version 7.0.1 only from the Version 7.0.1 MQ Explorer. If you connect the Version 8.0 MQ Explorer remotely to a Version 7.0.1 queue manager, no telemetry resources are displayed. You cannot connect a Version 8.0 MQ Explorer locally to a Version 7.0.1 queue manager on the same server.

IBM MQ Telemetry client libraries and SDK

To help you write messaging applications for MQTT networks, you can install and use a set of free client libraries. After you develop your applications, these applications and the client libraries are then deployed together to the appropriate system.

In previous versions of IBM MQ, the client libraries were supplied with the product, in the Client Software Development Kit (SDK). From IBM MQ Version 8.0, this SDK is no longer supplied as part of the product. Instead, the current version of the SDK is available as the free download IBM Messaging Telemetry Clients SupportPac.

To use the telemetry clients, download the IBM Messaging Telemetry Clients SupportPac, then install the clients in a directory of your own choosing. The sample applications and client libraries are in client-specific directories under `<CLIENTPACKDIR>/SDK/clients`, where `<CLIENTPACKDIR>` is the directory in which you uncompressed the client pack. Note that the example scripts included in IBM MQ assume that the clients are in the directories specified in Location of telemetry logs, error logs, and configuration files.

The IBM Messaging Telemetry Clients SupportPac provides you with the following resources:

- Sample MQTT client applications written in Java, in JavaScript, and in C.
- MQTT client libraries that support these client applications, and enable them to run on most platforms and devices, including Android devices and products from Apple.

Related information:

Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0

IBM MQ Telemetry

Telemetry use cases

Administering IBM MQ Telemetry

Developing applications for IBM MQ Telemetry

IBM MQ Telemetry reference

IBM MQ Telemetry troubleshooting

Planning your installation on Windows systems

This topic describes the different methods available to install IBM MQ on Windows systems and the different installation types.

If you are migrating from an earlier version of IBM MQ, see Planning IBM MQ migration to the latest version on UNIX and Linux platforms, Windows, and IBM i. To modify an existing installation, see “Modifying your installation” on page 293.

Interactive or Non-Interactive installation

IBM MQ for Windows is installed using the Microsoft Installer (MSI). You can use the Installation Launchpad to invoke MSI, this process is called an attended or interactive installation. Or, you can invoke MSI directly for a silent installation, without using the IBM MQ Installation Launchpad. This means that you can install IBM MQ on a system without interaction. This process is called unattended, silent, or non-interactive installation, and is useful for installing IBM MQ over a network on a remote system.

For a list of interactive and non-interactive features, see “IBM MQ features for Windows systems” on page 207.

Interactive installation

If you choose an interactive installation, before you install, you must decide what type of installation you require. Table 39 shows the installation types available, and the features that are installed with each option. For the prerequisites required for each feature, see IBM MQ System Requirements.

The installation types are:

- Typical installation
- Compact installation
- Custom Installation

You can also:

- Specify the installation location, name, and description.
- Have multiple installations on the same computer.

See “Choosing a primary installation” on page 184 for important information about these features, including whether to designate your installation as the *primary installation*.

Table 39. Features installed with each type of interactive installation

Installation type	Server Features installed	Client Features installed	Comments
Typical	<ul style="list-style-type: none">• Server• MQ Explorer• Development Toolkit• Java and .NET Messaging and Web Services	<ul style="list-style-type: none">• Windows Client• Development Toolkit• Java and .NET Messaging and Web Services	<p>The default option. Features are installed to default locations with a default installation name.</p> <p>Java and .NET Messaging and Web Services includes IBM MQ classes for .NET and support for the Microsoft Windows Communication Foundation (WCF) for use with Microsoft.NET 3.</p>
Compact	<ul style="list-style-type: none">• Server only	<ul style="list-style-type: none">• Windows Client only	<p>The feature is installed to the default location with a default installation name.</p>

Table 39. Features installed with each type of interactive installation (continued)

Installation type	Server Features installed	Client Features installed	Comments
Custom	<p>By default, the following features are preselected:</p> <ul style="list-style-type: none"> • Server • MQ Explorer • Development Toolkit • Java and .NET Messaging and Web Services <p>A custom installation can also install:</p> <ul style="list-style-type: none"> • Telemetry Service • Advanced Message Security • Managed File Transfer Service • Managed File Transfer Logger • Managed File Transfer Agent • Managed File Transfer Tools • Windows client 	<p>By default, the following features are preselected:</p> <ul style="list-style-type: none"> • Windows Client • Development Toolkit • Java and .NET Messaging and Web Services 	<p>A server custom installation can be used if you want to install the Windows client from within the server image.</p> <p>All the available features are listed and you can select which ones to install, and where to install them. You can also name and provide a description for the installation.</p> <p>Use a custom installation when you want to specify that the installation is primary.</p> <p>Java and .NET Messaging and Web Services includes IBM MQ classes for .NET and support for the Microsoft Windows Communication Foundation (WCF) for use with Microsoft.NET 3 or later.</p>

If Microsoft.NET is not installed before IBM MQ and you add it, rerun **setmqinst -i -n Installationname** if this is a primary installation.

The following table describes which level of .NET is required for which function:

Table 40. Required levels of Microsoft.NET

IBM MQ function	.NET version required
IBM MQ classes for .NET. For more information, see: Getting started with IBM MQ classes for .NET 2	.NET 2
<p>The IBM MQ custom channel for WCF. For more information, see Developing WCF applications with IBM MQ.</p> <p>To build the sample solution files, either the Microsoft.NET 3.5 SDK, or Microsoft Visual Studio 2008 is needed. For more information, see: Software requirements for the WCF custom channel for IBM MQ</p>	.NET framework 3.5 or later

For instructions on how to install IBM MQ on Windows systems, see Installing IBM MQ Server on Windows systems and “Installing an IBM MQ client on Windows systems” on page 337.

Non-interactive installation

If you choose a non-interactive installation the system on which you want to install must be able to access the IBM MQ image, or a copy of the files, and you must be able to access the system.

If you are running IBM WebSphere MQ Version 7.5 or later, with User Account Control (UAC) enabled, you must invoke the non-interactive installation from an elevated command prompt. Elevate a command

prompt by using a right-click to start the command prompt and choose **Run as administrator**. If you try to silently install from a non-elevated command prompt, the installation fails with an error of AMQ4353 in the installation log.

There are several ways to invoke MSI:

- Using the `msiexec` command with command-line parameters.
- Using the `msiexec` command with a parameter that specifies a response file. The response file contains the parameters that you normally supply during an interactive installation. See “Advanced installation using `msiexec`” on page 279.
- Use the `MQParms` command with command-line parameters, a parameter file, or both. The parameter file can contain many more parameters than a response file. See “Using the `MQParms` command” on page 287.

If the system belongs to a Windows domain you may need a special domain ID for the IBM MQ service, see “Security considerations when installing IBM MQ server on a Windows system” on page 219 for more information.

Clearing IBM MQ installation settings

When you install IBM MQ on Windows various values, such as the location of the data directory for IBM MQ, are stored in the registry.

In addition, the data directory contains configuration files that are read at install time. To provide a trouble free re-installation experience, these values and files persist even after the last IBM MQ installation has been removed from the machine.

This is designed to assist you, and

- Allows you to easily uninstall and reinstall
- Ensures that you do not lose any previously defined queue managers in the process.

However in some cases this feature can be an annoyance. For example, if you want to:

- Move the data directory
- Pick up the new default data directory for IBM MQ Version 8. See Windows: changes for IBM MQ Version 8. for further information.
- Install as if installing on a new machine, for example, for test purposes.
- Remove IBM MQ permanently.

To assist you in these situations, IBM MQ Version 8.0 onwards supplies a Windows command file, on the root directory of the installation media, called **ResetMQ.cmd**.

To run the command, enter the following:

```
ResetMQ.cmd [LOSEDATA] [NOPROMPT]
```

Attention: The parameters **LOSEDATA** and **NOPROMPT** are optional. If you supply either, or both, of these parameters, the following action results:

LOSEDATA

Existing queue managers become unusable. However, the data remains on disk.

NOPROMPT

Configuration information is permanently removed without further prompting.

You can run this command only after the last IBM MQ installation has been removed.

Important: You should use this script with caution. The command, even without specifying the optional parameter **LOSEDATA**, can irrecoverably remove queue manager configuration.

Related concepts:

“Security considerations when installing IBM MQ server on a Windows system”

Use this information to learn about user names and the security considerations when installing IBM MQ server on a Windows system.

Security considerations when installing IBM MQ server on a Windows system

Use this information to learn about user names and the security considerations when installing IBM MQ server on a Windows system.

- If you are installing IBM MQ on a Windows domain network running Active Directory Server, you probably need to obtain a special domain account from your domain administrator. For further information, and the details that the domain administrator needs to set up this special account, see *Configuring IBM MQ accounts*.
- When you are installing IBM MQ server on a Windows system you must have local administrator authority .
- In order to administer any queue manager on that system, or to run any of the IBM MQ control commands your user ID must belong to the *local mqm* or *Administrators* group . If the local mqm group does not exist on the local system, it is created automatically when IBM MQ is installed. A user ID can either belong to the local mqm group directly, or belong indirectly through the inclusion of global groups in the local mqm group.
- Windows versions with a User Account Control (UAC) feature restricts the actions users can perform on certain operating system facilities, even if they are members of the *Administrators* group. If your user ID is in the *Administrators* group but not the *mqm* group you must use an elevated command prompt to issue IBM MQ admin commands such as *crtmqm*, otherwise the error AMQ7077 is generated. To open an elevated command prompt, right-click the start menu item, or icon, for the command prompt, and select **Run as administrator**
- Some commands can be run without being a member of the *mqm* group (see *Authority to administer IBM MQ*).
- If you intend to administer queue managers on a remote system, your user ID must be authorized on the target system.
- As with other versions of Windows, the object authority manager (OAM) gives members of the *Administrators* group the authority to access all IBM MQ objects even when UAC is enabled.

Additional restrictions for installing on Windows

There are some additional points to consider when installing IBM WebSphere MQ Version 7.5 or later on Windows. First, Windows has some rules regarding the naming of objects created and used by IBM MQ. Second, you can set up logging during installation which assists you in troubleshooting any problems you might have with the installation.

Naming considerations

- Ensure that the machine name does not contain any spaces. IBM MQ does not support machine names that include spaces. If you install IBM MQ on such a machine, you cannot create any queue managers.
- For IBM MQ authorizations, names of user IDs and groups must be no longer than 64 characters (spaces are not allowed).
- An IBM MQ for Windows server does not support the connection of a Windows client if the client is running under a user ID that contains the @ character, for example, *abc@d*. Similarly, the client user ID should not be the same as local group.
- A user account that is used to run the IBM MQ Windows service is set up by default during the installation process; the default user ID is *MUSR_MQADMIN*. This account is reserved for use by IBM MQ. Refer to *Configuring IBM MQ accounts*.

- When an IBM MQ client connects to a queue manager on the server, the username under which the client runs must not be same as the domain or machine name. If the user has the same name as the domain or machine, the connection fails with return code 2035(MQRC_NOT_AUTHORIZED).

Digital signatures

The IBM MQ programs and installation image are digitally signed on Windows to confirm that they are genuine and unmodified. From IBM MQ Version 8.0 the SHA-256 with RSA algorithm is used to sign the IBM MQ product.

Logging

Logging is enabled by default from the Launchpad. You can also enable complete logging, for more information, see [How to enable Windows Installer logging](#)

Planning your installation on HP Integrity NonStop Server

This section describes what to do to prepare your system for installing IBM MQ client for HP Integrity NonStop Server .

Understanding multiple installations

IBM MQ client for HP Integrity NonStop Server can be installed more than once on an HP Integrity NonStop Server system. In addition, multiple different versions of IBM MQ can be installed on a single HP Integrity NonStop Server system, and be maintained independently. Each installation can be of any supported version of IBM MQ . There are no requirements for installations to be either the same, or different versions.

To install IBM MQ , you must specify two locations; one in the OSS file system, and one in the Guardian file system, which is used by the installer to store the results of the installation. These locations must not contain or overlap with any other IBM MQ installation. The locations must also be free of other files.

Each installation is independent and self-contained, with all data, such as configuration logs, or trace and program files located within the installation directory hierarchy. All commands and libraries use an embedded runtime search path (RPath) to ensure that they load their dependencies from the same installation.

As several installations might be present, each application must locate and load the IBM MQ client libraries from the correct installation.

- For native applications, an application that is linked with the IBM MQ MQIC.LIB installation library inherits the IBM MQ installation RPATH, and can run without environment variables. Environment variables in OSS, for example, `_RLD_LIB_PATH` or `DEFINES` in Guardian, are only required if you want to run the application using a different IBM MQ installation.
- For Java applications using the Java Messaging Service (JMS) API, the client Java archive (JAR) must be from the correct installation, and must be included on the class path. For more information, see [Environment variables used by IBM MQ classes for JMS](#).

Product packaging and delivery

IBM MQ client for HP Integrity NonStop Server is downloaded to the OSS file system as a single file.

The IBM MQ client for HP Integrity NonStop Server package file is a self-extracting archive (SFX) that contains an installer and all files that are required to create installations.

The SFX for IBM MQ client for HP Integrity NonStop Server has a file extension of .run. There is no concept of placed files. When run, the SFX creates a single installation, directly from the archive, into the OSS and Guardian file systems.

The SFX can be used to create as many installations of the IBM MQ client for HP Integrity NonStop Server as you require. No information about installations is retained in the SFX, and no tools are provided for extracting individual files from the SFX.

File system

Before you install the IBM MQ client for HP Integrity NonStop Server, make sure that the file system is set up correctly.

Review “Hardware and software requirements on HP Integrity NonStop Server systems” on page 231 to make sure that you understand the approximate amount of disk space in the OSS and Guardian file systems that is required for an installation. The OSS file set that is used for the installation requires enough free space for the installation files and the files you create in the installation. The Guardian volume that you use for installation does not require auditing.

Work with your systems administrator to verify the OSS file set and Guardian file system storage requirements, at least for an initial estimate of the storage. The best way to determine more precisely how much storage you would eventually need in production is to produce a prototype configuration and model the message storage requirements, scaling up as necessary for your production system.

OSS file system objects

For the OSS file system objects, this section concentrates on the differences between the HP Integrity NonStop Server installation, and the standard UNIX installation. Multiple independent installations are supported.

The opt and var trees must be present in a common root directory, which is selected at installation time. The opt tree contains files that do not change. For example, this tree contains program, library, dll, header files, and "static" data files. The var tree contains files that might change, and do hold status about the installation itself. Examples of files that this tree holds are configuration files, and log files.

Both the opt and var directories contain a single directory named mqm. The content of both trees is rooted in the opt/mqm and var/mqm directories.

This table displays a summary of the contents at the top-level of opt/mqm:

Table 41.

Directory	Purpose	Contents
bin	Contains the OSS programs and libraries for an installation	<ul style="list-style-type: none"> • G is a symbolic link file that locates the Guardian installation subvolume • amq* files, containing the product executables for the client • lib* files, containing product dll files • Files containing control commands, and other utilities and scripts

Table 41. (continued)

Directory	Purpose	Contents
inc	Contains the header files for building IBM MQ applications	<ul style="list-style-type: none"> • .h files, which are C language header files • .tal files which are pTAL header files • .cpy files which are COBOL copy file • cobcpy32 and cobcpy64 directories for the individual COBOL copy files
lib	Contains the import libraries needed to link applications	<ul style="list-style-type: none"> • G is a symbolic link file that locates the Guardian installation subvolume • amq* files, containing product dll files • iconv is a directory that contains data conversion tables • lib* files, which are product dll files • mqicb used for supplying to a CONSULT directive for compiling COBOL programs
license	Contains text versions of the IBM License for the IBM MQ client for HP Integrity NonStop Server product, which is translated into each supported national language	<ul style="list-style-type: none"> • Lic_.txt files, which are the individual national language translations of the license. • notices.txt is a file that contains any additional license terms from non-IBM software that is included with IBM MQ , if any
mq.id	Single file that contains information about the build level and install package	All contents in this directory might be used by IBM Support personnel.
msg	Contains globalization files for use by IBM MQ , in logging and displaying output in the supported national language translations	<p>The contents include:</p> <ul style="list-style-type: none"> • amq.cat globalization message catalog currently in use by the installation, created by the OSS utility "gencat" • amq.msg unprocessed globalization data that is used as input by gencat to create the catalog • Other minor files and directories that support the different translations

Table 41. (continued)

Directory	Purpose	Contents
samp	Contains sample code and executables to illustrate the use of IBM MQ	<ul style="list-style-type: none"> • *.cob sample COBOL language source files • *.c sample 'C' language source files • *.tal sample pTAL language source files • ccsid.new backup file of ccsid.tbl • ccsid.tbl file that contains a table of supported CCSIDs • *.ini sample configuration files • java directory that contains source for sample Java applications • jms directory that contains source for sample JMS applications • bin directory that contains executable versions of the samples • dlq directory that contains a source for the sample Dead Letter Queue Handler • preconnect directory that contains source for preconnect exit

For more information about the samples that are provided with IBM MQ client for HP Integrity NonStop Server , see Samples for IBM MQ client for HP Integrity NonStop Server.

This table displays a summary of the contents at the top-level of var/mqm:

Table 42.

Directory	Purpose	Contents
conv	Contains data conversion files	Binary data that supports the data conversion function for IBM MQ
errors	Contains installation-wide error logs and FDC files	Standard content, for example: <ul style="list-style-type: none"> • AMQERR01.LOG - current system-wide error log file • AMQERR02.LOG - previous system-wide error log file • AMQERR03.LOG - oldest system-wide error log file • *.FDC FFST files
exits	Stores DLLs containing exit code that is loaded by the queue managers in the installation	This file is empty at installation
log	Contains log files for recording and controlling units of work	Standard content
mqs.ini	The installation configuration file	Standard content
qmgrs	Directory beneath the location where all of the queue manager directories are created	Standard content

Table 42. (continued)

Directory	Purpose	Contents
sockets	Directory tree that contains various queue manager control files	Standard content
trace	Defined location to which trace data is written by IBM MQ	Standard content

Guardian installation subvolume

The Guardian single installation subvolume contains both the programs and libraries needed at runtime.

This table shows the contents of the Guardian installation subvolume:

Table 43.

File	Description
AMQINST	Internal file that describes installation configuration
AMQS*	Samples that are built for Guardian
B*SAMP	Sample build files for the various supported languages
CMQ*	Header files for the various supported languages, where files ending in: <ul style="list-style-type: none"> • h are C headers • T are pTAL headers • L are COBOL headers
MQ*	Product libraries
MQS*C	Sample C language source files
MQS*T	Sample pTAL language source files
MQS*L	Sample COBOL language source files

Control commands are also included, for a list, see HP Integrity NonStop Server client commands.





Checking requirements

Before you install IBM MQ, you must check for the latest information and system requirements.

About this task

A summary of the tasks that you must complete to check system requirements are listed here with links to further information.

Procedure

1. Check that you have the latest information, including information on hardware and software requirements. See “Finding product requirements and updated support information” on page 225.
2. Check that your systems meet the initial hardware and software requirements on your platform:
 -   “Hardware and software requirements on UNIX and Linux systems” on page 225
 -  “Hardware and software requirements on Windows systems” on page 227
 -  “Hardware and software requirements on IBM i systems” on page 229

The supported hardware and software environments are occasionally updated. See the IBM MQ System Requirements website for the latest information.

3. Check that your systems have sufficient disk space for the installation. See Disk space requirements.
4. Check that you have the correct licenses. See IBM MQ license information.

What to do next

When you have completed these tasks, you are ready to start preparing your systems for installation. For the next steps in installing IBM MQ, see “Preparing the system” on page 233.

Related concepts:

“Installing IBM MQ” on page 249

The topics in this section provide instructions on how to install IBM MQ.

“Uninstalling” on page 389

The topics in this section provide instructions on how to uninstall components.

Related information:

IBM MQ maintenance tasks

Finding product requirements and updated support information

Access the latest information for IBM MQ.

IBM MQ System Requirements website

For the latest details of hardware and software requirements on all supported platforms, see the IBM MQ System Requirements website. From IBM MQ Version 8.0, you can use the Software Product Compatibility Reports (SPCR) tool to find information on supported operating systems, system requirements, prerequisites, and optional supported software. For more information about the SPCR tool and links to reports for each supported platform, see the System Requirements for IBM MQ Version 8.0 web page.

Product readme file

The product readme file includes information about last minute changes and known problems and workarounds. The latest version is available on the product readmes web page. Always check to see that you have the latest copy.

Support information

The IBM MQ support web page is regularly updated with the latest product support information. For example, if you are migrating from an earlier version, look under the heading *Solve a problem* for the document *Problems and solutions when migrating*.

Related concepts:

“Installing IBM MQ” on page 249

The topics in this section provide instructions on how to install IBM MQ.

 “Installing IBM MQ for z/OS” on page 409

Use this topic to install the IBM MQ for z/OS product on your system.

Related information:

IBM MQ maintenance tasks

Troubleshooting and support

Hardware and software requirements on UNIX and Linux systems

Before you install IBM MQ, check that your system meets the hardware and operating system software requirements for the particular components you intend to install.

Hardware and software requirements are set out at System requirements for IBM MQ.

IBM MQ does not support host names that contain spaces. If you install IBM MQ on a system with a host name that contains spaces, you are unable to create any queue managers.

Java Message Service and SOAP transport

If you want to use Java Message Service and SOAP support, you need an IBM Java 7 SDK and Runtime Environment Version 7.0 or later.

For development, a JDK is required, and a JRE is required for running. The JRE does not need to be the JRE installed with IBM MQ, but has to be one from the supported list.

For a list of supported JDKs, see System requirements for IBM MQ.

On Linux : Apache Axis V1.4 provides support for SOAP and is shipped on the server DVD, but not installed.

For further information about SOAP with IBM MQ , see IBM MQ transport for SOAP.

On HP-UX : To run a 64-bit or 32-bit JVM use the -d64 or -d32 parameters on the command line when running a Java application to ensure the correct JVM is used.

On Linux: On the Power platform, the 32-bit and 64-bit JDKs are typically installed to different locations, for example, the 32-bit JDK is located in /opt/IBMJava2-ppc-50 and the 64-bit JDK is located in /opt/IBMJava2-ppc64-50. Ensure that the PATH variable is correctly set for your applications that use Java. To use the Postcard application described in "Verify the installation using the Postcard application" on page 370, you must use a 32-bit JDK.

On Solaris : The 32-bit and 64-bit JDKs are typically installed to the same directory. To run a 64-bit JVM use the -d64 or -d32 parameters on the command line when running a Java application to ensure the correct JVM is used.

You can check the version installed using the following command:

```
java -version
```

Secure Sockets Layer (SSL)

If you want to use the SSL support, you need the IBM Global Security Kit (GSKit) V8 package. This package is supplied with IBM MQ as one of the components available for installation.

HP-UX

To use SSL, IBM MQ clients on HP-UX must be built using POSIX threads.

Linux

Installing the g++ version runtime support

If you intend to run SSL channels then you must have the g++ runtime libraries installed. The GNU g++ libraries are called libgcc_s.so and libstdc++.so.6. On RPM based systems these are installed as part of the libgcc and libstdc++ software packages.

The version of these libraries installed must be compatible with g++ version 3.4.

See System requirements for IBM MQ for further details on the required packages for SSL support.

On 64 bit platforms, install both the 32 bit and the 64 bit versions of the package so that 32 bit and 64 bit processes can both use SSL functions.

MQ Explorer requirements

Linux MQ Explorer can be installed either as part of the product installation, or from the stand-alone MQ Explorer support pack MS0T. See MQ Explorer Requirements for the minimum requirements that your system needs, if you want to use the MQ Explorer .

Note that MQ Explorer is available for use only with IBM MQ for Linux, x86 and x86-64 platforms.

Unicode support on AIX

If you need to convert data to and from Unicode on your system, you must install the following file sets:

bos.iconv.ucs.com Unicode converters for AIX sets
bos.iconv.ucs.ebcdic Unicode converters for EBCDIC sets
bos.iconv.ucs.pc Unicode converters for PC sets

Solaris 11 operating system

If you are installing on the Solaris 11 operating system, ensure that the IPS package (package/svr4) that supports **pkgadd** and equivalent utilities is installed.

Related concepts:

 “Hardware and software requirements on IBM i systems” on page 229

Check that the server environment meets the prerequisites for installing IBM MQ for IBM i. Check the product readme files and install missing prerequisite software supplied on the server CD.

“Hardware and software requirements on Windows systems”

Check that the server environment meets the prerequisites for installing IBM MQ for Windows and install any prerequisite software that is missing from your system from the server DVD.

Related tasks:

“Checking requirements” on page 224

Before you install IBM MQ, you must check for the latest information and system requirements.

Hardware and software requirements on Windows systems

Check that the server environment meets the prerequisites for installing IBM MQ for Windows and install any prerequisite software that is missing from your system from the server DVD.

Before you install IBM MQ, you must check that your system meets the hardware and software requirements. For the latest details of hardware and software requirements on all supported platforms, see IBM MQ System Requirements.

You must also review the product readme file, which includes information about last-minute changes and known problems and workarounds. The latest version is always on the product readmes web page.

Storage requirements for IBM MQ server

The storage requirements depend on which components you install, and how much working space you need. The storage requirements also depend on the number of queues that you use, the number and size of the messages on the queues, and whether the messages are persistent. You also require archiving capacity on disk, tape, or other media. For more information, see IBM MQ System Requirements.

Disk storage is also required:

- Prerequisite software
- Optional software
- Your application programs

Requirements for MQ Explorer

MQ Explorer can be installed either as part of the product installation, or from the stand-alone MQ Explorer support pack MS0T.

- The product version is available for Windows x86_64.

- The support pack version is available for Windows x86 and x86_64.

The requirements for installing MQ Explorer as part of the product installation, and not as the stand-alone MQ Explorer support pack MS0T, include:

- A 64-bit (x86_64) processor
- 64-bit Windows operating system

If you are not running 64-bit Windows, you can install the 32-bit version of the MS0T support pack.

For further information about Windows requirements, see MQ Explorer Requirements and the following web pages:

- Windows 7 system requirements
- Windows 8 system requirements

Installation directories used for Windows operating systems

The 64-bit IBM MQ server or client, by default, installs its program directories into the 64-bit installation location: C:\Program Files\IBM\WebSphere MQ.

The 32-bit client version of IBM MQ, by default, installs its program directories to the 32-bit installation location: C:\Program Files (x86)\WebSphere MQ.

The default data directory that is used by IBM MQ has changed in Version 8.0 to C:\ProgramData\IBM\MQ. This change affects both servers and clients, in 32 and 64 bits. However, if there has been a previous installation of IBM MQ on the machine on which you are installing, the new installation continues to use the existing data directory location. For more information, see Program and data directory locations.

Installing prerequisite software

To install the prerequisite software that is provided on the IBM MQ Server DVD (which does not include service packs or web browsers), choose one of the following options:

- Use the IBM MQ installation procedure.


When you install using the IBM MQ Server DVD, there is a **Software Prerequisites** option in the IBM MQ Installation Launchpad window. You can use this option to check what prerequisite software is already installed and what is missing, and then install any missing software.

- Use Windows Explorer:
 1. Use Windows Explorer to select the Prereqs folder on the IBM MQ Server DVD.
 2. Select the folder for the software item to be installed.
 3. Start the installation program.

Related concepts:

“Hardware and software requirements on UNIX and Linux systems” on page 225

Before you install IBM MQ, check that your system meets the hardware and operating system software requirements for the particular components you intend to install.

 “Hardware and software requirements on IBM i systems”

Check that the server environment meets the prerequisites for installing IBM MQ for IBM i. Check the product readme files and install missing prerequisite software supplied on the server CD.

Related tasks:

“Checking requirements” on page 224

Before you install IBM MQ, you must check for the latest information and system requirements.

Related information:

MQ Explorer Requirements

Hardware and software requirements on IBM i systems



Check that the server environment meets the prerequisites for installing IBM MQ for IBM i. Check the product readme files and install missing prerequisite software supplied on the server CD.

Before installation, you must check that your system meets the hardware and software requirements set out in the IBM MQ system requirements page at <http://www.ibm.com/software/integration/wmq/requirements/>. You must also review the release notes file, which are on the product CD in the \Readmes folder for each national language, and check the READADD.txt file for any changes made between translation and the manufacturing of the installation CD. READADD.txt is found in the root directory of the server installation CD.

During installation, the release notes file is copied to the IBM MQ program files folder (default /QIBM/ProdData/mqm).

Storage requirements for IBM MQ server

The storage requirements for IBM MQ for IBM i depend on which components you install, and how much working space you need. The storage requirements also depend on the number of queues that you use, the number and size of the messages on the queues, and whether the messages are persistent. You also require archiving capacity on disk, tape, or other media. For more information, see the IBM MQ system requirements page at, <http://www.ibm.com/software/integration/wmq/requirements/>.

Disk storage is also required:

- Prerequisite software
- Optional software
- Your application programs

Installing prerequisite software

To install the prerequisite software provided on the IBM MQ Server CD (which does not include service packs or web browsers), do one of the following:

- Use the IBM MQ installation procedure.

When you install using the IBM MQ Server CD, there is a **Software Prerequisites** option in the IBM MQ Installation Launchpad window. You can use this option to check what prerequisite software is already installed and which is missing, and to install any missing software.

Using TLS Version 1.2

TLS version 1.2 is the latest version of the Secure Sockets Layer (SSL) protocol. The core System TLS v1.2 functionality is included in IBM i 7.1 Technology Refresh 6 (TR6). To enable and use the new protocols, program temporary fixes (PTFs) from multiple areas of the operating system are also required.

Provided DCM (5770SS1 option 34) is installed on your system, requesting and applying SI48659 obtains all of the enablement PTFs.

System value changes

The new support is installed, but dormant in System SSL after applying SI48659.

In order to activate the new protocols for System SSL, use Change System Value (CHGSYSVAL) to modify The QSSLPCL system value.

Change the default value of *OPSYS to:

- *TLSV1.2
- *TLSV1.1
- *TLSV1
- *SSLV3

If QSSLPCL is set to something other than *OPSYS, add *TLSV1.2 and *TLSV1.1 to the existing setting.

Prerequisite PTFs for multiple certificate support

You are not limited to a single certificate for SSL/TLS channels. To use multiple certificates on IBM i platforms, you must install the following program temporary fixes (PTFs):

- MF57749
- MF57889
- SI52214
- MF58003

See Digital certificate labels: understanding the requirements for details about how to select certificates by using certificate labels.

Related concepts:

  “Hardware and software requirements on UNIX and Linux systems” on page 225

Before you install IBM MQ, check that your system meets the hardware and operating system software requirements for the particular components you intend to install.

 “Hardware and software requirements on Windows systems” on page 227

Check that the server environment meets the prerequisites for installing IBM MQ for Windows and install any prerequisite software that is missing from your system from the server DVD.

“License requirements” on page 233

You must have purchased sufficient licenses for your installation. The details of the license agreement is stored on your system at installation time so that you can read it at any time. IBM MQ supports IBM License Metric Tool (ILMT).

Related tasks:

“Checking requirements” on page 224

Before you install IBM MQ, you must check for the latest information and system requirements.

Hardware and software requirements on HP Integrity NonStop Server systems

Check that the server environment meets the prerequisites for installing the IBM MQ client for HP Integrity NonStop Server. Check the product readme files and install missing prerequisite software supplied on the server CD.

Hardware

The IBM MQ client for HP Integrity NonStop Server typically requires certain hardware specifications to run:

- HP Integrity NonStop Server H and J series
- Two or more processors
- At least 1 GB, and ideally 4 GB of memory per processor
- 500 MB of free disk space in the Guardian and OSS file systems

Operating system

Two operating systems are supported by the IBM MQ client for HP Integrity NonStop Server:

- HP Integrity NonStop Server running H06.24 or later NonStop OS
- HP Integrity NonStop BladeSystem running J06.13 or later NonStop OS

You must be running one of these operating systems to install the IBM MQ client for HP Integrity NonStop Server.

Other software requirements

IBM MQ client for HP Integrity NonStop Server has some additional software requirements:

- The operating system software, Open System Services (OSS), must be active, with file systems and a local sockets subsystem that is configured and running.
- Safeguard must be active.
- If two-phase commit transaction support is required, then TMF must be active and Pathway must be configured and available. The connected queue manager must be at IBM WebSphere MQ Version 7.1 or later.
- If the Java Message Service (JMS) API is required, then HP Integrity NonStop Server for Java V6 must be available.
- You might require compatible compilers, linkers, and maybe other tools for the C, C++, COBOL, JMS, or pTAL languages if you want to build and use applications.

File system requirements

In the selected installation root directory, in the OSS file system, an installation creates:

- opt - a directory tree that contains the "static" files for an installation in OSS.
- var - a directory tree that contains the "variable" files for an installation in OSS.

An installation also creates a single subvolume in the Guardian file system, which is selected during installation.

Related concepts:

“Finding product requirements and updated support information” on page 225

Access the latest information for IBM MQ.

“License requirements” on page 233

You must have purchased sufficient licenses for your installation. The details of the license agreement is stored on your system at installation time so that you can read it at any time. IBM MQ supports IBM License Metric Tool (ILMT).

Related information:

  Disk space requirements

Verifying system software prerequisites:

Use the HP Integrity NonStop Server TACL utility, `SYSINFO`, to verify the base OS level of the HP Integrity NonStop Server.

Procedure

From a TACL command prompt, enter `SYSINFO`.

Results

The system information is displayed as shown in the following example:

```
SYSINFO - T9268H01 - (01 OCT 2004) SYSTEM \NODE1 Date 05 Nov 2010, 11:56:51  
Copyright 2003 Hewlett-Packard Development Company, L.P.
```

```
      System name      \NODE1  
EXPAND node number    025  
      Current SYSnn    SYS00  
      System number    nnnnnn  
Software release ID    J06.10.00
```

In this example, the base OS level is J06.10.00.

What to do next

Compare the base OS level with the “Hardware and software requirements on HP Integrity NonStop Server systems” on page 231. Verify any other HP Integrity NonStop Server software prerequisites or recommendations identified in the documentation or the product README; for example, SPRs to particular products.

License requirements

You must have purchased sufficient licenses for your installation. The details of the license agreement is stored on your system at installation time so that you can read it at any time. IBM MQ supports IBM License Metric Tool (ILMT).

Important: Ensure that your enterprise has the correct license, or licenses, for the components that you are going to install. See IBM MQ license information for more details.

License files

At installation, the license agreement files are copied into the `/licenses` directory under the `MQ_INSTALLATION_PATH`. You can read them at any time.



 On IBM i, you can use the `WRKSWAGR` command to view the software licenses.


ILMT


ILMT automatically detects IBM MQ, if you are using it, and checks with it each time a queue manager is started. You do not need to take any further action. You can install ILMT before or after IBM MQ.

The automatic detection applies to both the IBM MQ server and IBM MQ Java products.

Related concepts:

  “Hardware and software requirements on UNIX and Linux systems” on page 225
Before you install IBM MQ, check that your system meets the hardware and operating system software requirements for the particular components you intend to install.

 “Hardware and software requirements on IBM i systems” on page 229
Check that the server environment meets the prerequisites for installing IBM MQ for IBM i. Check the product readme files and install missing prerequisite software supplied on the server CD.

 “Hardware and software requirements on Windows systems” on page 227
Check that the server environment meets the prerequisites for installing IBM MQ for Windows and install any prerequisite software that is missing from your system from the server DVD.

Related tasks:

“Checking requirements” on page 224
Before you install IBM MQ, you must check for the latest information and system requirements.





Preparing the system







On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

About this task

The tasks that you perform to prepare your systems for installation are listed here. Complete the appropriate tasks for your platform before installing.

Procedure

-   On UNIX and Linux systems, set up the user and group. See “Setting up the user and group on UNIX and Linux systems” on page 234
-   On UNIX and Linux, create file systems. See “Creating file systems on UNIX and Linux systems” on page 236
- Configure additional settings for your platform:

-  “Operating System configuration and tuning for IBM MQ on AIX systems” on page 239
-  “Operating System configuration and tuning for IBM MQ on HP-UX systems” on page 239
-  “Operating System configuration and tuning for IBM MQ on IBM i” on page 247
-  “Operating system configuration and tuning for IBM MQ on Linux systems” on page 242
-  “Operating System configuration and tuning for IBM MQ on Solaris systems” on page 245
-  “Additional restrictions for installing on Windows” on page 219

What to do next

When you have completed the tasks to prepare the system, you are ready to start installing IBM MQ. To install a server, see “Installing an IBM MQ server” on page 254. To install a client, see “Installing an IBM MQ client” on page 319.

Related concepts:

 “Installing IBM MQ for z/OS” on page 409

Use this topic to install the IBM MQ for z/OS product on your system.

Related information:

Planning

Migrating and upgrading IBM MQ

IBM MQ maintenance tasks

Setting up the user and group on UNIX and Linux systems

On UNIX and Linux systems, IBM MQ requires a user ID of the name `mqm`, with a primary group of `mqm`. The `mqm` user ID owns the directories and files that contain the resources associated with the product.

Creating the user ID and group on UNIX and Linux systems

Set the primary group of the `mqm` user to the group `mqm`.

If you are installing IBM MQ on multiple systems you might want to ensure each UID and GID of `mqm` has the same value on all systems. If you are planning to configure multi-instance queue managers, it is essential the UID and GID are the same from system to system. It is also important to have the same UID and GID values in virtualization scenarios.

AIX

You can use the System Management Interface Tool (`smit`), for which you require root authority.

1. To create the `mqm` group, display the required window using this sequence:

```
Security & Users
Groups
Add a Group
```

Set the group name field to `mqm`.

2. To create the user `mqm`, display the required window using this sequence:

```
Security & Users
Users
Add a User
```

Set the user name field to `mqm`.

3. To add a password to the new user ID, display the required window using this sequence:

Security & Users
Passwords
Change a User's Password

Set the password as required.

HP-UX

The user ID value for user `mqm` must be less than 60,000 to avoid problems with the maintenance update process.

You can use the System Management Homepage (SMH), or the **groupadd** and **useradd** commands to work with user IDs.

Linux RPM creates the `mqm` user ID and group ID as part of the installation procedure if they do not exist.

If you have special requirements for these IDs (for example they need to have the same values as other machines you are using, or your users and group ID are centrally managed) you should create the IDs before running the installation procedure, using the **groupadd** and **useradd** commands to set the UID and GID the same on each machine.

Note: The only IBM MQ requirement, is that the `mqm` user should have the `mqm` group as its primary group.

Solaris

The user ID value for user `mqm` must be less than 262,143 to avoid problems with the maintenance update process.

Create the IDs using the **groupadd** and **useradd** commands to set the UID and GID the same on each machine.

Adding existing user IDs to the group on UNIX and Linux systems

If you want to run administration commands, for example **crtmqm** (create queue manager) or **strmqm** (start queue manager), your user ID must be a member of the `mqm` group. This user ID must not be longer than 12 characters.

Users do not need `mqm` group authority to run applications that use the queue manager; it is needed only for the administration commands.

AIX

You can use **smit** to add an existing user ID to the `mqm` group. Display the required menu using this sequence:

```
Security & Users  
Users  
Change / Show Characteristics of a User
```

Type the name of the user in the **User Name** field and press **Enter**. Add `mqm` to the **Group SET** field, which is a comma-separated list of the groups to which the user belongs. Users do not need to have their primary group set to `mqm`. If `mqm` is in their set of groups, they can use the administration commands.

Log files created by IBM MQ Telemetry service

The **umask** setting of the user ID that creates a queue manager will determine the permissions of the Telemetry log files generated for that queue manager. Even though the ownership of the log files will be set to `mqm`.

Related concepts:

“Creating file systems on UNIX and Linux systems”

Before installing IBM MQ Version 8.0 , you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

“Operating System configuration and tuning for IBM MQ on AIX systems” on page 239

“Operating System configuration and tuning for IBM MQ on HP-UX systems” on page 239

Before you install IBM MQ on an HP-UX system, you must check that the kernel is configured correctly.

“Operating system configuration and tuning for IBM MQ on Linux systems” on page 242

Use this topic to when configuring IBM MQ on Linux systems.

Related information:

“Operating System configuration and tuning for IBM MQ on Solaris systems” on page 245

Configure Solaris systems with the resource limits required by IBM MQ.

Setting up the user and group on HP Integrity NonStop Server

The administrator user ID must be used to administer the IBM MQ client for HP Integrity NonStop Server.

Ensure that you have access to an IBM MQ client for HP Integrity NonStop Server user ID in the user group called MQM. The MQM group must be created before the client can be installed. All user IDs that are used to install the client must have MQM as their primary group. If this user group does not exist, or you do not have access to such a user, contact your systems administrator.

Creating file systems on UNIX and Linux systems

Before installing IBM MQ Version 8.0 , you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

Determining the size of a server installations file system

To determine the size of the `/var/mqm` file system for a server installation, consider:

- The maximum number of messages in the system at one time.
- Contingency for message buildups, if there is a system problem.
- The average size of the message data, plus 500 bytes for the message header.
- The number of queues.
- The size of log files and error messages.
- The amount of trace that is written to the `/var/mqm/trace` directory.

Storage requirements for IBM MQ also depend on which components you install, and how much working space you need. For more details, see [Disk space requirements](#).

Creating a file system for the working data

Before you install IBM MQ, create and mount a file system called `/var/mqm` which is owned by the user `mqm` in the group `mqm` ; see “Setting up the user and group on UNIX and Linux systems” on page 234. This file system is used by all installations of IBM MQ on a system. If possible, use a partition strategy with a separate volume for the IBM MQ data. This means that other system activity is not affected if a large amount of IBM MQ work builds up. Configure the directory permissions to permit the `mqm` user to have full control, for example, file mode 755. These permissions will then be updated during the IBM MQ installation to match the permissions required by the queue manager.

Creating separate file systems for errors and logs

You can also create separate file systems for your log data (/var/mqm/log) and error files (/var/mqm/errors). If possible, place these directories on different physical disks from the queue manager data (/var/mqm/qmgrs) and from each other.

If you create separate file systems the /var/mqm/errors directory can be NFS mounted. However, if you choose to NFS-mount /var/mqm/errors, the error logs might be lost if the network fails.

You can protect the stability of your queue manager by having separate file systems for:

- /var/mqm/errors
- /var/mqm/trace
- /var/mqm/qmgrs
- /var/mqm/log

In the case of /var/mqm/errors, it is rare that this directory receives large quantities of data. But it is sometimes seen, particularly if there is a severe system problem leading to IBM MQ writing a lot of diagnostic information in to .FDC files. In the case of /var/mqm/trace, files are only written here when you use **strmqtrc** to start tracing IBM MQ.

You can obtain better performance of normal IBM MQ operations (for example, syncpoints, MQPUT, MQGET of persistent messages) by placing the following on separate disks:

- /var/mqm/qmgrs
- /var/mqm/log

In the rare event that you need to trace an IBM MQ system for problem determination, you can reduce performance impact by placing the /var/mqm/trace file system on a separate disk.

If you are creating separate file systems, allow a minimum of 30 MB of storage for /var/mqm, 100 MB of storage for /var/mqm/log, and 10 MB of storage for /var/mqm/errors. The 100 MB minimum allowance of storage for /var/mqm/log is the absolute minimum required for a single queue manager and is not a recommended value. The size of a file system must be scaled according to the number of queue managers that you intend to use, the number of pages per log file, and the number of log files per queue manager.

If you want to use individual queues that hold more than 2 GB of data, you must enable /var/mqm to use large files.

For more information about file systems, see File system support.

The size of the log file depends on the log settings that you use. The minimum sizes are for circular logging using the default settings. For more information about log sizes, see Calculating the size of the log.

Linux and Solaris

For a client installation, the file system can be mounted on a remote network device, for example NFS.

If you are performing both a client and a server installation, the requirements of the server installation take precedence over the requirements of the client installation.

Allow 15 MB as a minimum for an IBM MQ client.

A new sample IBM MQ MQI client configuration file is created in the var/mqm directory, by the client package, during installation, but only if this file does not exist. This file contains the ClientExitPath stanza. An example mqclient.ini file is shown in Configuring a client using a configuration file.

If you are using a common configuration file for multiple clients, either in the IBM MQ installation directory or in another location using the MQCLNTCF environment variable, you must grant read access to all user identifiers under which the IBM MQ client applications run. If, for any reason, the file cannot be read the failure is traced, and the search logic continues as if the file had not existed.

Related concepts:

“Setting up the user and group on UNIX and Linux systems” on page 234

On UNIX and Linux systems, IBM MQ requires a user ID of the name mqm, with a primary group of mqm. The mqm user ID owns the directories and files that contain the resources associated with the product.


“Operating System configuration and tuning for IBM MQ on AIX systems” on page 239

“Operating System configuration and tuning for IBM MQ on HP-UX systems” on page 239

Before you install IBM MQ on an HP-UX system, you must check that the kernel is configured correctly.

“Operating system configuration and tuning for IBM MQ on Linux systems” on page 242

Use this topic to when configuring IBM MQ on Linux systems.

 “Operating System configuration and tuning for IBM MQ on IBM i” on page 247

Before installing IBM MQ for IBM i, there are several system values which need to be checked using the DSPSYSVAL command. If necessary, reset the values using the CHGSYSVAL command.

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

Related information:


“Operating System configuration and tuning for IBM MQ on Solaris systems” on page 245

Configure Solaris systems with the resource limits required by IBM MQ.

Operating System configuration and tuning for IBM MQ on UNIX and Linux systems

, and IBM i

Some UNIX and Linux and Linux systems require you to make additional settings.

- “Operating System configuration and tuning for IBM MQ on AIX systems” on page 239
- “Operating System configuration and tuning for IBM MQ on HP-UX systems” on page 239
- “Operating system configuration and tuning for IBM MQ on Linux systems” on page 242
- “Operating System configuration and tuning for IBM MQ on Solaris systems” on page 245
-  “Operating System configuration and tuning for IBM MQ on IBM i” on page 247

Operating System configuration and tuning for IBM MQ on AIX systems

File descriptors

When running a multi-threaded process such as the agent process, you might reach the soft limit for file descriptors. This limit gives you the IBM MQ reason code MQRC_UNEXPECTED_ERROR (2195) and, if there are enough file descriptors, an IBM MQ FFST file.

To avoid this problem, increase the process limit for the number of file descriptors. You must alter the `nfiles` attribute in `/etc/security/limits` to 10,000 for the `mqm` user ID, or in the default stanza. To alter the number of file descriptors do these steps:

1. In a command prompt, check the maximum number of file descriptors available to a process running as `mqm`:

```
lsuser -a nfiles mqm
```
2. Set the value to at least 10240:

```
chuser nfiles_hard=10240 mqm  
chuser nfiles=10240 mqm
```

System Resource Limits

Set the system resource limit for data segment and stack segment to unlimited using the following commands in a command prompt:

```
ulimit -d unlimited  
ulimit -s unlimited
```

Attention: For an `mqm` user ID other than root, the value `unlimited` might not be permitted.

For more information on configuring your system, see *How to configure UNIX and Linux systems for IBM MQ*.

You can check your system configuration using the `mqconfig` command.

Related concepts:

“Setting up the user and group on UNIX and Linux systems” on page 234

On UNIX and Linux systems, IBM MQ requires a user ID of the name `mqm`, with a primary group of `mqm`. The `mqm` user ID owns the directories and files that contain the resources associated with the product.

“Creating file systems on UNIX and Linux systems” on page 236

Before installing IBM MQ Version 8.0, you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

Operating System configuration and tuning for IBM MQ on HP-UX systems

Before you install IBM MQ on an HP-UX system, you must check that the kernel is configured correctly.

Kernel configuration

It is possible that the default kernel configuration is not adequate because IBM MQ uses semaphores and shared memory.

Before installation, review the configuration of the machine and increase the values if necessary. Consider using the values of the tunable kernel parameters given in Table 44. These values might need to be increased if you obtain any First Failure Support Technology™ (FFST) records.

Note:

1. Semaphore and swap usage do not vary significantly with message rate or message persistence.
2. IBM MQ queue managers are independent of each other. Therefore system tunable kernel parameters, for example `shmmni`, `semmni`, `semmns`, and `semmnu` need to allow for the number of queue managers in the system.

See the HP-UX documentation for information about changing these values.

Table 44. Minimum tunable kernel parameters values

Name	Value	Increase	Description
<code>shmmax</code>	268435456	No	Maximum size of a shared-memory segment (bytes)
<code>shmseg</code>	1024	No	Maximum number of shared memory segments per process
<code>shmmni</code>	1024	Yes	Maximum number of shared memory segments
<code>semaem</code>	128	No	Maximum undo value for a semaphore for a single process
<code>semvmx</code>	32767	No	Maximum value of a semaphore
<code>semmns</code>	4096	Yes	Maximum number of semaphores
<code>semmni</code>	128	Yes	Maximum number of semaphore sets
<code>semmnu</code>	16384	Yes	Maximum number of process having semaphore operations that can be undone
<code>semume</code>	32	No	Maximum number of semaphore undo operations per process
<code>max_thread_proc</code>	66	No	Maximum number of threads in a process
<code>maxfiles</code>	10000	No	Maximum number of file handles per process (soft limit)
<code>maxfiles_lim</code>	10000	No	Maximum number of file handles per process (hard limit)

Notes:

- These values are sufficient to run two moderate sized queue managers on the system. If you intend to run more than two queue managers, or the queue managers are to process a significant workload, you might need to increase the values displayed as *Yes* in the *Increase* column.
- You must restart the system after you change any of the tunable kernel parameters.

System resource limits

You can set global limits for the size of process data segments and the size of process stack segments for the whole system. These limits are set by altering the tunable kernel parameters.

The tunable kernel parameters are:

Parameter	What it controls	Consider minimum value
maxdsiz	Maximum size of the data segment for 32-bit processes	1073741824
maxdsiz_64bit	Maximum size of the data segment for 64-bit processes	1073741824
maxssiz	Maximum size of the stack segment for 32-bit processes	8388608
maxssiz_64bit	Maximum size of the stack segment for 64-bit processes	8388608

If other software on the same machine needs higher values, then the operation of IBM MQ is not adversely affected if those higher values are used.

For the full documentation for these parameters see the HP-UX product documentation.

To apply the settings to an HP-UX 11i system which has the System Administration Manager (SAM) utility, you can use SAM to achieve the following steps:

- Select and alter the parameters
- Process the new kernel
- Apply the changes and restart the system

Other releases of HP-UX might provide different facilities to set the tunable kernel parameters. Consult your HP-UX product documentation for the relevant information.

The `ulimit` shell command

On a per-shell basis, the available limits can be tuned down from the values stored for the “System resource limits” on page 240 preceding parameters. Use the `ulimit` shell command to tune the values of the parameters with a combination of the following switches:

Switch	Meaning
-H	The hard limit
-S	The soft limit
-d	The data segments size
-s	The stack segment size

Verifying that the kernel settings are applied

You can verify that the resource limits have not been lowered by a `ulimit` command and that the queue manager has the correct limits. To verify the limits, go to the shell from which the queue manager is started and enter the following command:

```
ulimit -Ha
ulimit -Sa
```

Among the console output you see:

```
data(kbytes) 1048576
stack(kbytes) 8192
```

If the lowered numbers are returned, then a `ulimit` command has been issued in the current shell to reduce the limits. Consult with your system administrator to resolve the issue.

You can check your system configuration using the `mqconfig` command.

For more information on configuring your system, see How to configure UNIX and Linux systems for IBM MQ.

Related concepts:

“Setting up the user and group on UNIX and Linux systems” on page 234

On UNIX and Linux systems, IBM MQ requires a user ID of the name `mqm`, with a primary group of `mqm`. The `mqm` user ID owns the directories and files that contain the resources associated with the product.

“Creating file systems on UNIX and Linux systems” on page 236

Before installing IBM MQ Version 8.0 , you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

Operating system configuration and tuning for IBM MQ on Linux systems

Use this topic to when configuring IBM MQ on Linux systems.

Attention: The information in this topic applies only if the queue manager is started by the `mqm` user ID.

If any other user ID starts the queue manager, ensure that the **NOFILE** and **NPROC** entries, shown for `mqm`, are duplicated for that user ID.

Shell interpreter

Ensure that `/bin/sh` shell is a valid shell interpreter compatible with the Bourne shell, otherwise the post-installation configuration of IBM MQ does not complete successfully. If the shell was not installed using RPM, you might see a prerequisites failure of `/bin/sh` shell when you try to install IBM MQ . The failure is because the RPM tables do not recognize that a valid shell interpreter is installed. If the failure occurs, you can reinstall the `/bin/sh` shell by using RPM, or specify the RPM option `--nodeps` to disable dependency checking during installation of IBM MQ .

Note: The `--dbpath` option is not supported when installing IBM MQ on Linux.

System V IPC kernel configuration

IBM MQ uses System V IPC resources, in particular shared memory. However, a limited number of semaphores are also used.

The minimum configuration for IBM MQ for these resources is as follows:

Table 45. Minimum tunable kernel parameters values

Name	Kernel-name	Value	Increase	Description
<code>shmmni</code>	<code>kernel.shmmni</code>	4096	Yes	Maximum number of shared memory segments
<code>shmmax</code>	<code>kernel.shmmax</code>	268435456	No	Maximum size of a shared-memory segment (bytes)
<code>shmall</code>	<code>kernel.shmall</code>	2097152	Yes	Maximum amount of shared memory (pages)

Table 45. Minimum tunable kernel parameters values (continued)

Name	Kernel-name	Value	Increase	Description
semmsl	kernel.sem	32	No	Maximum amount of semaphores permitted per set
semms	kernel.sem	4096	Yes	Maximum number of semaphores
semopm	kernel.sem	32	No	Maximum number of operations in single operations
semnmi	kernel.sem	128	Yes	Maximum number of semaphore sets
thrmax	kernel.threads-max	1000	Yes	Maximum number of threads
pidmax	kernel.pid_max	12000	Yes	Maximum number of process identifiers

Notes:

1. These values are sufficient to run two moderate sized queue managers on the system. If you intend to run more than two queue managers, or the queue managers are to process a significant workload, you might need to increase the values displayed as Yes in the Increase column.
2. The kernel.sem values are contained within a single kernel parameter containing the four values in order.

To view the current value of the parameter log on, as a user with root authority, and type:

```
sysctl <Kernel-name>
```

To add or alter these values, log on as a user with root authority. Open the file /etc/sysctl.conf with a text editor, then add or change the following entries to your chosen values:

```
kernel.shmmni = 4096
kernel.shmall = 2097152
kernel.shmmax = 268435456
kernel.sem = 32 4096 32 128
```

Then save and close the file.

To load these **sysctl** values immediately, enter the following command `sysctl -p`.

If you do not issue the `sysctl -p` command, the new values are loaded when the system is rebooted.

By default the Linux kernel has a maximum process identifier, that can also be used with threads, and might limit the allowed number of threads.

The operating system reports when the system lacks the necessary resources to create another thread, or the system-imposed limit on the total number of threads in a process {PTHREAD_THREADS_MAX} would be exceeded.

For more information on kernel.threads-max and kernel.pid-max, see Resource shortage in IBM MQ queue manager when running a large number of clients

TCP/IP configuration

If you want to use **keepalive** for IBM MQ channels, you can configure the operation of the KEEPALIVE using the kernel parameters:

```
net.ipv4.tcp_keepalive_intvl
net.ipv4.tcp_keepalive_probes
net.ipv4.tcp_keepalive_time
```

See Using the TCP/IP SO_KEEPALIVE option for further information.

To view the current value of the parameter `log on`, as a user with root authority, and type `sysctl <Kernel-name>`.

To add or alter these values, log on as a user with root authority. Open the file `/etc/sysctl.conf` with a text editor, then add or change the following entries to your chosen values.

To load these **sysctl** values immediately, enter the following command `sysctl -p`.

If you do not issue the `sysctl -p` command, the new values are loaded when the system is rebooted.

Maximum open files

The maximum number of open file-handles in the system is controlled by the parameter **fs.file-max**

The minimum value for this parameter for a system with two moderate sized queue managers is 524288.

If you intend to run more than two queue managers, or the queue managers are to process a significant workload, you might need to increase this value.

To view the current value of a parameter, log on as a user with root authority, and type `sysctl fs.file-max`.

To add or alter these values, log on as a user with root authority. Open the file `/etc/sysctl.conf` with a text editor, then add or change the following entry to your chosen value:

```
fs.file-max = 524288
```

Then save and close the file.

To load these **sysctl** values immediately, enter the following command `sysctl -p`.

If you do not issue the `sysctl -p` command, the new values are loaded when the system is rebooted.

If you are using a pluggable security module such as PAM (Pluggable Authentication Module), ensure that this module does not unduly restrict the number of open files for the `mqm` user. To report the maximum number of open file descriptors per process for the `mqm` user, login as the `mqm` user and enter the following values:

```
ulimit -n
```

For a standard IBM MQ queue manager, set the *nofile* value for the `mqm` user to 10240 or more. To set the maximum number of open file descriptors for processes running under the `mqm` user, add the following information to the `/etc/security/limits.conf` file:

```
mqm      hard nofile    10240
mqm      soft nofile    10240
```

Maximum processes

A running IBM MQ queue manager consists of a number of thread programs. Each connected application increases the number of threads running in the queue manager processes. It is normal for an operating system to limit the maximum number of processes that a user runs. The limit prevents operating system failures due to an individual user or subsystem creating too many processes. You must ensure that the maximum number of processes that the `mqm` user is allowed to run is sufficient. The number of processes must include the number of channels and applications that connect to the queue manager.

The following calculation is useful when determining the number of processes for the `mqm` user:

```
nproc = 2048 + clientConnections * 4 + qmgrChannels * 4 +
        localBindingConnections
```

where:

- *clientConnections* is the maximum number of connections from clients on other machines connecting to queue managers on this machine.
- *qmgrChannels* is the maximum number of running channels (as opposed to channel definitions) to other queue managers. This includes cluster channels, sender/receiver channels, and so on.
- *localBindingConnections* does not include application threads.

The following assumptions are made in this algorithm:

- 2048 is a large enough contingency to cover the queue manager threads. This might need to be increased if a lot of other applications are running.
- When setting *nproc*, take into account the maximum number of applications, connections, channels and queue managers that might be run on the machine in the future.
- This algorithm takes a pessimistic view and the actual *nproc* needed might be slightly lower for later versions of IBM MQ and fastpath channels.
- **V8.0.0.7** In Linux, each thread is implemented as a light-weight process (LWP) and each LWP is counted as one process against *nproc*.

You can use the `PAM_limits` security module to control the number of processes that users run. You can configure the maximum number of processes for the `mqm` user as follows:

```
mqm      hard  nproc    4096
mqm      soft  nproc    4096
```

For more details on how to configure the `PAM_limits` security module type, enter the following command:

```
man limits.conf
```

You can check your system configuration using the `mqconfig` command.

For more information on configuring your system, see *How to configure UNIX and Linux systems for IBM MQ*.

Related concepts:

“Setting up the user and group on UNIX and Linux systems” on page 234

On UNIX and Linux systems, IBM MQ requires a user ID of the name `mqm`, with a primary group of `mqm`. The `mqm` user ID owns the directories and files that contain the resources associated with the product.

“Creating file systems on UNIX and Linux systems” on page 236

Before installing IBM MQ Version 8.0, you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

Operating System configuration and tuning for IBM MQ on Solaris systems

Configure Solaris systems with the resource limits required by IBM MQ.

IBM MQ uses semaphores, shared memory, and file descriptors, and it is probable that the default resource limits are not adequate.

For further information on **maxusers**, and other process-sizing parameters, on Solaris 10, see Process sizing parameters.

To set new default limits for all users in the *mqm* group, set up a project for the *mqm* group in each zone.

To find out if you already have a project for the *mqm* group, log in as root and enter the following command:

```
projects -l
```

If you do not already have a *group.mqm* project defined, enter the following command:

```
projadd -c "WebSphere MQ default settings"  
-K "process.max-file-descriptor=(basic,10000,deny)"  
-K "project.max-shm-memory=(priv,4GB,deny)"  
-K "project.max-shm-ids=(priv,1024,deny)"  
-K "project.max-sem-ids=(priv,128,deny)" group.mqm
```

If a project called *group.mqm* is listed, review the attributes for that project. The attributes must include the following minimum values:

```
process.max-file-descriptor=(basic,10000,deny)  
project.max-sem-ids=(priv,128,deny)  
project.max-shm-ids=(priv,1024,deny)  
project.max-shm-memory=(priv,4294967296,deny)
```

If you need to change any of these values, enter the following command:

```
projmod -s -K "process.max-file-descriptor=(basic,10000,deny)"  
-K "project.max-shm-memory=(priv,4GB,deny)"  
-K "project.max-shm-ids=(priv,1024,deny)"  
-K "project.max-sem-ids=(priv,128,deny)" group.mqm
```

Note that you can omit any attributes from this command that are already correct.

For example, to change only the number of file descriptors, enter the following command:

```
projmod -s -K "process.max-file-descriptor=(basic,10000,deny)" group.mqm
```

(To set only the limits for starting the queue manager under the *mqm* user, login as *mqm* and enter the command `projects`. The first listed project is likely to be `default`, and so you can use `default` instead of `group.mqm`, with the `projmod` command.)

To ensure that the attributes for the project *group.mqm* are used by a user session when running IBM MQ, make sure that the primary group of that user ID is *mqm*. In the examples in this topic, the *group.mqm* project ID will be used.

For further information on how projects are associated with user sessions, see *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones* for your release of Solaris.

You can check your system configuration using the `mqconfig` command.

For more information on configuring your system, see *How to configure UNIX and Linux systems for IBM MQ*.

Related concepts:

“Setting up the user and group on UNIX and Linux systems” on page 234

On UNIX and Linux systems, IBM MQ requires a user ID of the name `mqm`, with a primary group of `mqm`. The `mqm` user ID owns the directories and files that contain the resources associated with the product.

“Creating file systems on UNIX and Linux systems” on page 236

Before installing IBM MQ Version 8.0, you might need to create file systems for both the product code and working data to be stored. There are minimum storage requirements for these file systems. The default installation directory for the product code can be changed at installation time, but the working data location cannot be changed.

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

Operating System configuration and tuning for IBM MQ on IBM i



Before installing IBM MQ for IBM i, there are several system values which need to be checked using the `DSPSYSVAL` command. If necessary, reset the values using the `CHGSYSVAL` command.

Check the following values and change if required:

QCCSID

Every message has a coded-character set identifier (CCSID) in its header. The CCSID tag identifies the code page and character set of the source.

A queue manager obtains its CCSID from the job that created it. If the job CCSID is not a valid value in the range 1-65534, the queue manager uses the default CCSID value (65535) instead. You can change the CCSID used by the IBM MQ queue manager by using the `CL` command **CHGMQM**.

Note: The CCSID must be either single-byte character set (SBCS), or mixed, that is SBCS and DBCS. It must not be DBCS only.

QSYSLIBL

Ensure that `QSYS2` is included in the list of libraries that make up the system part of the library list. IBM MQ uses programs in this library for data conversion and SNA LU 6.2 communication.

Note: Do not have `QMQM` as part of the system or user portion of the library list.

QALWOBJRST

Ensure that the `QALWOBJRST` system value is set to `*ALL` or `*ALWPGMADP` before you install MQ. If it is set to `*NONE`, installation fails.

After installation, reset `QALWOBJRST` to its original value to maintain system security.

QSHRMEMCTL

Ensure that the `QSHRMEMCTL` system value is set to 1 (Allowed).

A value of 1 is used in environments where pointers can be shared amongst programs between different jobs.

IBM MQ requires this setting to use the shared memory APIs `shmat` and `shmget` and to share its pointers across jobs.

If it is not set correctly, initialization of IBM MQ fails with system return code "3401" (Permission denied), and commands such as `CRTMQM`, `STRMQM`, `ENDMQM`, `TRCMQM` fail.

QFRCCVNRST

Ensure that the QFRCCVNRST system value is set to 0 (Restore all objects without conversion), or 1 (Objects with validation errors are converted), before you install MQ. If it is not set, installation fails.

QMLTTHDACN

Optionally set this to control the generation of messages into joblogs. Set QMLTTHDACN to 2 to get messages generated in a joblog; set it to 1 to avoid the messages. For example, the message CPD000D is an informational message that is generated when a command that is not thread-safe is issued from a multi-threaded application. Setting QMLTTHDACN to 1 avoids the message.

Related concepts:

“Hardware and software requirements on IBM i systems” on page 229

Check that the server environment meets the prerequisites for installing IBM MQ for IBM i. Check the product readme files and install missing prerequisite software supplied on the server CD.

“License requirements” on page 233

You must have purchased sufficient licenses for your installation. The details of the license agreement is stored on your system at installation time so that you can read it at any time. IBM MQ supports IBM License Metric Tool (ILMT).

Related tasks:

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

“Installing IBM MQ server on IBM i” on page 303

Install IBM MQ for IBM i by installing the IBM MQ server in its primary language, installing samples and installing additional languages.



Installing IBM MQ

The topics in this section provide instructions on how to install IBM MQ.

See “Finding product requirements and updated support information” on page 225 for details about how to check that you have access to the most recent information available.

If product fixes or updates are made available, see IBM MQ maintenance tasks for information about how to apply these changes.

To prepare for installation and to install the IBM MQ components, complete the following tasks:

- “Planning your installation” on page 181
- “Checking requirements” on page 224
- “Preparing the system” on page 233
- “Installing an IBM MQ server” on page 254
- “Installing an IBM MQ client” on page 319
- “Installing IBM MQ Telemetry” on page 214
-  “Installing IBM MQ Java messaging and web services for IBM i” on page 359
-  “Installing IBM MQ for z/OS” on page 409

Related concepts:

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

“Verifying a client installation” on page 374

You can verify that your IBM MQ MQI client installation completed successfully and that the communication link is working.

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

Related tasks:

Installing IBM MQ Advanced Message Security

Install and uninstall IBM MQ Advanced Message Security component.

Related information:

Migrating and upgrading IBM MQ

IBM MQ Managed File Transfer product options

Installation using Electronic Software Download

You can perform an installation of IBM MQ from an installation image downloaded from IBM.

Before you start installing, consider how you want to use IBM MQ and review the general Planning section.

Go to the Passport Advantage[®] and Passport Advantage Express web site for further information on how you:

- Acquire new IBM software licenses.
- Renew Software Subscription and Support and Fixed Term Licenses.
- Buy and renew technical support for some Selected Open Source and other non-warranted applications.

- Subscribe to IBM SaaS offerings and acquire IBM Appliances.

Passport Advantage is designed for larger enterprises and enterprises with multiple sites.

Passport Advantage Express is designed for smaller enterprises and single-site enterprises.

Redistributable clients

► V8.0.0.4

The IBM MQ redistributable client is a collection of runtime files that are provided in a .zip or .tar file that can be redistributed to third parties under redistributable license terms, which provides a simple way of distributing applications and the runtime files that they require in a single package.

What are the IBM MQ distributable clients?

From IBM MQ Version 8.0.0, Fix Pack 4, native redistributable client runtime libraries are provided for Linux x86-64 and Windows 64-bit platforms to make it simple to distribute both applications and the required IBM MQ runtime libraries. A third package, which is not platform-specific, contains the runtime files that are required for the Java/JMS applications, including the IBM MQ resource adapter for JMS applications that are running under an application server.

The redistributable client that is supplied with IBM MQ, is also a *non-installed* and *relocatable* image. Maintenance of a *redistributable, non-installed* image, is achieved through replacement; that is, you download newer versions of the runtime components when they are shipped.

A *redistributable* client implies distributing the required runtime with an application both inside and outside of your environment.

A *relocatable* client implies putting the files somewhere else other than a fixed default location. For example, instead of installing into `/opt/` installing into `/usr/local`.

A *non-installed* client implies that you are not required to lay down client files, and that these files can be copied as required.

The IBM IPLA license agreement is being extended for IBM MQ to enable you to download a number of additional runtime files from Fix Central.

Supported languages

You can use the files that are contained in the redistributable images to run the following client applications:

- C
- C++
- COBOL
- Java
- Java JMS
- Both fully managed and unmanaged .NET

Limitations

GSKit objects

No new GSKit objects are being shipped. Only the runtime files are shipped, both in a regular installation and with the redistributable client.

IBM JREs

No IBM JREs are being provided with the redistributable client.

If you want to run Java/JMS applications, you must provide your own runtime environment. Your JRE, that applications run under, must meet the current SOE requirements and are bound by any restrictions or limitations that apply.

Developing applications

All other files that support the development and distribution of applications (including copybooks, header files, and sample source code) are not included in the redistributable client and are not licensed for redistribution.

If you need to develop IBM MQ applications, you still need to perform a traditional installation so that you obtain the SDK files that are required to build client applications.

Windows C runtime libraries

You might have these libraries on your machine already, but if you do not, you need to download and install the following Microsoft C/C++ runtime libraries:

- Microsoft Visual C++ Redistributable 2008
- Microsoft Visual C++ Redistributable 2012

The download links for the redistributable downloads for each of these libraries can be found at The latest supported Visual C++ downloads .

Related concepts:

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Choosing an installation location” on page 192

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the *MQ_INSTALLATION_PATH*.

“Installation considerations for redistributable clients”

The Linux x86-64 image is shipped in a LinuxX64.tar.gz file, and the Windows 64-bit image is shipped in a Win64.zip file.

“.NET application runtime - Windows only” on page 253

Considerations when using the .NET application.

Installation considerations for redistributable clients

► V8.0.0.4

The Linux x86-64 image is shipped in a LinuxX64.tar.gz file, and the Windows 64-bit image is shipped in a Win64.zip file.

File names

The archive or .zip file names describe the file contents and equivalent maintenance levels. For example, in IBM MQ Version 8.0.0, Fix Pack 4 the client images are available under the following file names:

Linux x86-64

8.0.0.4-WS-MQC-Redist-LinuxX64.tar.gz

Windows

8.0.0.4-WS-MQC-Redist-Win64.zip

Java - all platforms

8.0.0.4-WS-MQC-Redist-Java.zip

Choosing the runtime files to distribute with an application

A script file named **genmqpkg** is provided by the redistributable client under the `bin` directory.

You can use the **genmqpkg** script to generate a smaller subset of files that are tailored to the needs of the application, for which the files are intended to be distributed.

You are asked a series of interactive Yes or No questions to determine the runtime requirements for an IBM MQ application.

Finally, **genmqpkg** asks you to supply a new target directory, where the script duplicates the required directories and files.

Important: IBM support is only able to provide assistance with the full, unmodified set of files contained within the redistributable client packages.

Other considerations

There is a minor change to the default data path of a non-installed client, and now the path is, on:

Linux x86-64

```
$HOME/IBM/MQ/data
```

Windows

```
%HOMEDRIVE%%HOMEPATH%\IBM\MQ\data
```

A redistributable client runtime co-exists with a full IBM MQ client or server installation, provided that they are installed in different locations.

Important: Unpacking a redistributable image into the same location as a full IBM MQ installation is not supported.

On Linux the `ccsid.tbl` used to define the supported CCSID conversions is traditionally expected to be found in the `UserData` directory structure, along with error logs, trace files, and so on.

The `UserData` directory structure is populated by unpacking the redistributable client, and so, if the file is not found in its usual location, the redistributable client falls back to locate the file in the `/lib` subdirectory of the installation.

Classpath changes

The classpath used by **dspmqver**, **setmqenv**, and **crtmqenv** commands, add the `com.ibm.mq.allclient.jar` to the environment, immediately following the `com.ibm.mq.jar` and `com.ibm.mqjms.jar`.

An example of **dspmqver** output from the redistributable client on Linux:

```
Name:          IBM MQ
Version:       8.0.0.4
Level:         p800-804-L150909
BuildType:     IKAP - (Production)
Platform:      IBM MQ for Linux (x86-64 platform)
Mode:          64-bit
O/S:           Linux 2.6.32.59-0.7-default
InstName:      MQNI08000004
InstDesc:      IBM MQ V8.0.0.4 (Redistributable)
Primary:       No
InstPath:      /Development/johndoe/unzip/unpack
DataPath:      /u/johndoe/IBM/MQ/data
MaxCmdLevel:  802
```

An example of **dspmqver** output from the redistributable client on Windows:

Name: IBM MQ
Version: 8.0.0.4
Level: p800-804-L150909
BuildType: IKAP - (Production)
Platform: IBM MQ for Windows (x64 platform)
Mode: 64-bit
O/S: Windows 7 Professional x64 Edition, Build 7601: SP1
InstName: MQNI08000004
InstDesc: IBM MQ V8.0.0.4 (Redistributable)
Primary: No
InstPath: C:\Users\johndoe\Desktop\Redist
DataPath: C:\Users\johndoe\IBM\MQ\data
MaxCmdLevel: 802

Related concepts:

“Redistributable clients” on page 250

The IBM MQ redistributable client is a collection of runtime files that are provided in a .zip or .tar file that can be redistributed to third parties under redistributable license terms, which provides a simple way of distributing applications and the runtime files that they require in a single package.

“.NET application runtime - Windows only”

Considerations when using the .NET application.

.NET application runtime - Windows only



Considerations when using the .NET application.

The runtime DLL files laid down in the *redistributable* images on Windows for .NET applications are normally registered with the global assembly cache (GAC) by a user with system administrator privileges, when installing the primary installation. However, this severely limits the benefits of redistribution.

The *redistributable* package on the Windows platform does not provide any tooling to register DLLs with the GAC, so .NET applications have to locate the appropriate assemblies by other means. There are two options that work in this situation.

Probing

After checking the GAC, the .NET runtime attempts to locate required assemblies through probing. The first location checked is the application base, which is the root location where the application is being run. See the information on *How the Runtime Locates Assemblies* on the Microsoft Web site for more information.

Note that when using this approach, the maintenance level of the assemblies used when building the .NET application must match those used at runtime – for example an application built at IBM MQ Version 8.0.0, Fix Pack 4 must be run with the IBM MQ Version 8.0.0, Fix Pack 4 redistributable client runtime.

Using this approach, a .NET application placed in the \bin directory alongside the IBM MQ assemblies picks up assemblies from a primary IBM MQ installation (if one exists), falling back to the redistributable copies.

1. Compile the .NET application under a full IBM MQ installation, that is `csc \t:exe \r:System.dll \r:amqmdnet.dll \lib: \out:nmqwrld.exe nmqwrld.cs`.
2. Copy the exe file in the redistributable client zip file into the \bin directory.

DEVPATH environment variable

An alternative, that allows your application to be built, distributed, unzipped and run as previously, is to use DEVPATH to locate the required assemblies. Unlike with the probing approach, this option overrides any matching assemblies from the GAC. However it is for this reason that Microsoft discourages its use in a production environment.

This approach can be effective where there is a possibility that a full IBM MQ installation is installed on the client. However, there is a good reason to always use the redistributable assemblies.

1. Compile the .NET application under a full IBM MQ installation, that is `csc \t:exe \r:System.dll \r:amqmdnet.dll \lib: \out:nmqwrld.exe nmqwrld.cs`
2. Copy the exe file into, or alongside, the redistributable client zip file.
3. In the same directory as the exe, create an application configuration file with the name of the exe file suffixed by `.config`, that is `nmqwrld.exe.config` with the following contents:

```
<configuration>
  <runtime>
    <developmentMode developerInstallation="true" />
  </runtime>
</configuration>
```

4. Call `setmqenv -s` and set the `DEVPATH` environment variable to specify the `\bin` directory from the redistributable image before running the application, that is:

```
set DEVPATH=%MQ_INSTALLATION_PATH%\bin
```

Starting and stopping trace for the .NET redistributable managed client

You generate trace for the .NET redistributable managed client in the same way as for the stand-alone .NET client. For more information, see [Using the stand-alone IBM MQ .NET client](#).

More information on .NET

For more information on .NET, see [Writing and deploying IBM MQ .NET programs](#).

Related concepts:

“Redistributable clients” on page 250

The IBM MQ redistributable client is a collection of runtime files that are provided in a `.zip` or `.tar` file that can be redistributed to third parties under redistributable license terms, which provides a simple way of distributing applications and the runtime files that they require in a single package.

“Installation considerations for redistributable clients” on page 251

The Linux x86-64 image is shipped in a `LinuxX64.tar.gz` file, and the Windows 64-bit image is shipped in a `Win64.zip` file.

Installing an IBM MQ server

After preparing your system for installation you can install IBM MQ by following the appropriate instructions for your platform. After installation, you might want to verify your installation to check that installation has been successful.

Before you begin

Make sure that you have prepared your system. See [Preparing the system](#).


About this task

It is possible to have both a server and a client installation on the same machine; for instructions on how to do this see, “Installing an IBM MQ client” on page 319.

IBM MQ Telemetry is installed as part of the IBM MQ Server installation. It must be selected as part of a custom installation. For more information, see “Installing IBM MQ Telemetry” on page 214

Procedure

To begin the installation procedure, select the appropriate platform and installation method.

- Interactive installation
 1. “Installing IBM MQ server on AIX” on page 256
 2. “Installing IBM MQ server on HP-UX” on page 259
 3. “Installing IBM MQ server on Linux” on page 263
 4. “Installing on Linux Ubuntu or Linux on POWER Systems - Little Endian ” on page 267
 5. “Installing IBM MQ server on Solaris” on page 272
 6. “Installing IBM MQ server on Windows” on page 277
 7.  “Installing IBM MQ server on IBM i” on page 303
- Non-interactive installation
 1. “Non-interactive installation of IBM MQ server on AIX” on page 258
 2. “Non-interactive installation of IBM MQ server on HP-UX” on page 261
 3. “Non-interactive installation of IBM MQ server on Linux” on page 266
 4. “Non-interactive installation of IBM MQ server on Solaris” on page 275
 5. “Non-interactive installation of IBM MQ server on IBM i” on page 306
 6. “Advanced installation using msixexec” on page 279

Related concepts:

“Planning your installation” on page 181

Before you install IBM MQ, you must choose which components to install and where to install them. You must also make some platform-specific choices.

“Installing an IBM MQ client” on page 319

“Installing IBM MQ Telemetry” on page 214

From IBM WebSphere MQ Version 7.1, IBM MQ Telemetry is a component of the main IBM MQ product, and is no longer a separate plug-in. You can choose to install IBM MQ Telemetry when you first install IBM MQ, or when you modify an existing IBM MQ installation. To get a set of client libraries that help you write messaging applications for telemetry, download the free IBM Messaging Telemetry Clients SupportPac.

Related tasks:

“Displaying messages in your national language on UNIX and Linux systems” on page 317

To display messages from a different national language message catalog, you must install the appropriate catalog and set the **LANG** environment variable.

“Checking requirements” on page 224

Before you install IBM MQ, you must check for the latest information and system requirements.

“Preparing the system” on page 233

On some operating systems, you might have to complete several tasks before you install IBM MQ depending on your installation platform. You might also want to complete other tasks, depending on your installation intentions.

“Converting a trial license on UNIX, Linux, and Windows” on page 316

Convert a trial license to a full license without reinstalling IBM MQ.

Installing IBM MQ server on AIX

Before you begin

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- IBM MQ can be installed into System Workload Partitions (WPARs) with both shared and private file systems. For installation into private file systems, IBM MQ can be installed directly into the System WPAR using the procedure outlined in this topic. For installation into shared file systems, see Installing IBM MQ in AIX Workload Partitions. There are some limitations for shared /usr file systems:
 - The **dspmqinst** and **dspmqver** commands might report the primary installation incorrectly when compared with the symbolic links in /usr/bin. To synchronize the reporting of the primary installation in a System WPAR and the global environment, run **setmqinst** with the **-i** or **-x** parameter, on the individual zones.
 - You cannot change the primary installation within a WPAR. You must change the primary installation through the global environment, which has appropriate write access to /usr/bin.

Note: During installation to a non-default location, ATTENTION messages that relate to **errupdate** or **trcupdate** are produced. These messages are not errors. However, AIX system trace for IBM MQ is not supported for installations in a non-default location, and IBM MQ trace must be used for problem determination.

- If you install a copy of IBM MQ server for AIX using Electronic Software Download, obtained from Passport Advantage, you need to:
 - Uncompress the tar file, by using the following command:

```
uncompress WS_MQ_V8.0_TRIAL_FOR_AIX_ML.tar.Z
```
 - Extract the installation files from the tar file, by using the following command:

```
tar -xvf WS_MQ_V8.0_TRIAL_FOR_AIX_ML.tar
```
 - Use the installation tools, **installp** or **smit**, to install the IBM MQ server for AIX.

Tip: If you find that the Function keys do not work in SMIT, try pressing Esc and the Function key number to emulate the required Function key.

About this task

IBM MQ is supplied as a set of filesets that are installed using the standard AIX installation tools. The procedure uses the system management interface tool (SMIT), but you can choose to use **installp**, **geninstall** or the web-based System Manager. You can select which components you want to install. The components and file sets are listed in “Choosing what to install” on page 194.

This procedure installs IBM MQ into the default location of /usr/mqm.

If you want to install IBM MQ in any one of the following situations:

- As the first installation on your system using **installp**
- As the first installation on your system, and you are installing the product to a location that is not the default
- Alongside an existing installation

use the procedure described in “Non-interactive installation of IBM MQ server on AIX” on page 258.

If you want to carry out a side-by-side installation, alongside an existing installation of IBM MQ in the default location, the existing installation must be IBM WebSphere MQ Version 7.0.1.6, or later.

You must install the second version of the product in a location that is not the default, using **installp** (see “Non-interactive installation of IBM MQ server on AIX” on page 258).

If you want to carry out a single stage migration, refer to UNIX, Linux , and Windows: Single-stage migration from Version 7.0.1 , or later, to the latest version

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Select the required **smit** window using the following sequence:
Software Installation and Maintenance
Install and Update Software
Install and Update from ALL Available Software
4. List the software in the **SOFTWARE to install** field:
 - a. Enter **.**
 - b. Press **F4**
5. Select the filesets to install from the list. Ensure that you include the appropriate message catalog if you require messages in a language different from the language specified by the locale selected on your system. Enter **ALL** to install all applicable filesets.
6. View the license agreement:
 - a. Change **Preview new LICENSE agreements?** to **yes**
 - b. Press **Enter**
7. Accept the license agreements and install IBM MQ:
 - a. Change **ACCEPT new license agreements?** to **yes**
 - b. Change **Preview new LICENSE agreements?** to **no**
 - c. Press **Enter**

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.
You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).
- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ . For more information, see [setmqenv](#) and [crtmqenv](#).
- If you want to confirm that the installation was successful, you can verify your installation. See [“Verifying a server installation”](#) on page 363, for more information.

Related concepts:

“Choosing an installation location” on page 192

You can install IBM MQ into the default location. Alternatively, you can install into a custom location during the installation process. The location where IBM MQ is installed is known as the `MQ_INSTALLATION_PATH`.

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Non-interactive installation of IBM MQ server on AIX”

Install IBM MQ server from the command line using the AIX `installp` command.

“Uninstalling IBM MQ on AIX” on page 389

On AIX, you can uninstall the IBM MQ server or client using the System Management Interface Tool (SMIT) or the `installp` command.

Related information:

setmqinst

Changing the primary installation

Non-interactive installation of IBM MQ server on AIX

Install IBM MQ server from the command line using the AIX `installp` command.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

Note: During installation, errors relating to `errupdate` or `trcupdate` might occur. This can be caused by installing to a non-default location, if so these errors can be safely ignored. However, native trace for IBM MQ is only supported when installed in the default location.

About this task

You can use this method to install to a non-default location, and can select which components you want to install. The components and filesets are listed in “Choosing what to install” on page 194.

Procedure

1. Log in as root, or switch to the superuser using the `su` command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the CD, a network location, or a local file system directory.
3. Install the product in one of the following ways:
 - Install the whole product in the default location:
`installp -acgXYd . all`
 - Install selected file sets in the default location:
`installp -acgXYd . list of file sets`
 - Install the whole product in a non-default location using the `-R` flag:
`installp -R USIL_Directory -acgXYd . all`
 - Install selected file sets in a non-default location using the `-R` flag:
`installp -R USIL_Directory -acgXYd . list of file sets`

where *USIL_Directory* is a directory which exists before the command is run; it must not contain any spaces or *usr/mqm*. IBM MQ is installed underneath the directory specified. For example, if */USIL1* is specified, the IBM MQ product files are located in */USIL1/usr/mqm*. This location is known as the *MQ_INSTALLATION_PATH*.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see *Changing the primary installation*.

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see *setmqenv* and *crtmqenv*.
- If you want to confirm that the installation was successful, you can verify your installation. See *“Verifying a server installation”* on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Installing IBM MQ server on AIX” on page 256

“Uninstalling IBM MQ on AIX” on page 389

On AIX, you can uninstall the IBM MQ server or client using the System Management Interface Tool (SMIT) or the **installp** command.

Related information:

setmqinst

Changing the primary installation

User Specified Installation Location (USIL)

Installing IBM MQ server on HP-UX

Before you begin

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in *“Preparing the system”* on page 233.
- If you install a copy of IBM MQ server for HP-UX using Electronic Software Download, obtained from Passport Advantage, you need to uncompress the *tar.gz* file, and extract the installation files from the tar file, by using the following command:

```
tar -xvf WS_MQ_V8.0_TRIAL_FOR_HP-UX_ML.tar
```

Important: You must use GNU tar (also known as *gtar*) to unpack any tar images.

About this task

This task describes the installation of a server, using the **swinstall** program to select which components you want to install. The components are listed in *“Choosing what to install”* on page 194.

Note: If you are using a screen reader, use the non-interactive installation option “Non-interactive installation of IBM MQ server on HP-UX” on page 261, so that you can accept the license without viewing it.

If you are installing IBM MQ from a depot that contains service update packages, read HP-UX: Applying maintenance level updates on IBM MQ before installing the service update packages.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Accept the license by running the `mqlicense` script:

```
./mqlicense.sh
```

The license is displayed. If you accept the license, you can continue the installation.

4. Start the interactive installation procedure by typing the following command,
`swinstall -s installation_file`

installation_file is the absolute path to the installation file. The path must begin with a / and end with the name of the installation file. The installation file has a file name extension of `.v11`. In the resulting menu screen, select **MQSERIES**.

a. If you do not want to install all IBM MQ components, open **MQSERIES**

- 1) Mark the components you want to install. The installer resolves dependencies automatically.
- 2) Review the information displayed by the installer.

5. Optional: To install IBM MQ to a non-default location, select **MQSERIES** from the lower part of the user interface, then select **Actions > Change Product Location**. The default installation location is `/opt/mqm`.

For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

Note: Ensure that you do not select **Actions > Change Target** by accident, they are not the same.

6. If this installation is not the first installation on the system, select **Options > Allow creation of multiple versions**
7. Select **Actions > Install**. The log file tells you if there are any problems that need fixing.
8. Fix any problems, and click **OK** to install. You are informed when the installation has finished.
9. If this installation is not the first installation on the system, you must enter the following command to configure IBM MQ.

Note: `MQ_INSTALLATION_PATH` is the path where you have just installed IBM MQ and the character defining the path is a lower case L.

```
swconfig -x allow_multiple_versions=true MQSERIES,l=MQ_INSTALLATION_PATH
```

If you do not enter this command, the `swlist` command reports the installation as installed instead of configured. You must not use IBM MQ unless the installation is configured.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see *Changing the primary installation*.

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see *setmqenv* and *crtmqenv*.
- If you want to confirm that the installation was successful, you can verify your installation. See *“Verifying a server installation”* on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Non-interactive installation of IBM MQ server on HP-UX”

You can perform a non-interactive installation of IBM MQ using the **swinstall** command. A non-interactive installation is also known as a silent, or unattended installation.

“Uninstalling IBM MQ on HP-UX” on page 391

On HP-UX, you can uninstall the IBM MQ server or client using the **swremove** command.

Related information:

setmqinst

Changing the primary installation

Non-interactive installation of IBM MQ server on HP-UX

You can perform a non-interactive installation of IBM MQ using the **swinstall** command. A non-interactive installation is also known as a silent, or unattended installation.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in *“Preparing the system”* on page 233.

About this task

This topic describes the non-interactive installation of a server, using the **swinstall** program to select which components you want to install. The components and are listed in *“Choosing what to install”* on page 194.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the CD, a network location, or a local file system directory.
3. Accept the IBM MQ license agreement without an interactive prompt by entering the following command:

```
./mqlicense.sh -accept
```
4. Install IBM MQ using the **swinstall** command:
 - a. If this installation is not the first installation on the system, you must add **-x allow_multiple_versions=true** to the **swinstall** command.
 - b. Add the names of the components to install as parameters of the **swinstall** command. The installer automatically resolves any dependencies.

- c. Optional: Identify the installation location by adding `,l=MQ_INSTALLATION_PATH` as a parameter of the **swinstall** command. For each installation, all of the IBM MQ components that you require must be installed in the same location. The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

For example, to install all IBM MQ components, in a non-default location, as the first installation, enter the following command:

```
swinstall -s /installation_file.v11 MQSERIES,l=/opt/customLocation
```

To perform a partial installation, providing a list of components, in the default location, as the second installation, enter the following command:

```
swinstall -x allow_multiple_versions=true -s /installation_file.v11  
MQSERIES.MQM-RUNTIME MQSERIES.MQM-BASE MQSERIES.MQM-SERVER
```

/installation_file.v11 is the absolute path to the installation file. The path must begin with a / and end with the name of the installation file. The installation file has the extension *.v11*.

5. If this installation is not the first installation on the system, you must enter the following command to configure the installation:

Note: *MQ_INSTALLATION_PATH* is the path where you have just installed IBM MQ and the character defining the path is a lower case L.

```
swconfig -x allow_multiple_versions=true MQSERIES,l=MQ_INSTALLATION_PATH
```

If you do not enter this command, the **swlist** command reports the installation as installed instead of configured. You must not use IBM MQ unless the installation is configured.

Example

The example shows the command to run a silent, full installation in the default location, using the alternative form of specifying the source depot using `-x source_directory=` instead of `-s`. Notice that all the language features are installed. Run a partial installation to install your chosen languages.

```
cd /downloads/WMQInstallFiles  
swinstall -v -x source_directory=$PWD/hpUxxxx.v11 MQSERIES
```

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- If you want to confirm that the installation was successful, you can verify your installation. See [“Verifying a server installation”](#) on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Installing IBM MQ server on HP-UX” on page 259

“Uninstalling IBM MQ on HP-UX” on page 391

On HP-UX, you can uninstall the IBM MQ server or client using the **swremove** command.

Related information:

setmqinst

Changing the primary installation

Installing IBM MQ server on Linux

You can install an IBM MQ server on a 64-bit Linux system.

Before you begin

- If you install a copy of IBM MQ server for Linux using Electronic Software Download, obtained from Passport Advantage, you need to uncompress the tar.gz file, and extract the installation files from the tar file, by using the following command:

```
tar -xvf WS_MQ_V8.0_TRIAL_FOR_Linux_ML.tar
```

Important: You must use GNU tar (also known as gtar) to unpack any tar images.

- **CAUTION:**

The instructions in this topic do not apply to Linux Ubuntu or Linux on POWER Systems - Little Endian . See “Installing on Linux Ubuntu or Linux on POWER Systems - Little Endian ” on page 267 for these platforms.

- Before you start the installation procedure, ensure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- If this installation is not the first installation on the system, you must ensure that the **crtmqpkg** command can write to a temporary location. By default, the **crtmqpkg** command will write to the /var/tmp directory. To use a different location, you can set the *TMPDIR* environment variable before you run the **crtmqpkg** command.
- To run the **crtmqpkg** command used in this task, you must have the **pax** command or **rpmbuild** installed.

Attention: **pax** and **rpmbuild** are not supplied as part of the product. You must obtain these from your Linux distribution supplier.

About this task

Install the server by using the RPM Package Manager installer to select the components you want to install. The components and package names are listed in “Installing on Linux Ubuntu or Linux on POWER Systems - Little Endian ” on page 267.

Attention: If you install the packages using the wildcard character, that is, using the command `rpm -ivh MQ*.rpm`, you should install the packages in the following order:

- MQSeriesRuntime
- MQSeriesJRE
- MQSeriesJava
- MQSeriesServer
- MQSeriesFTBase
- MQSeriesFTAgent
- MQSeriesFTService
- MQSeriesFTLogger
- MQSeriesFTTools
- MQSeriesAMQP
- MQSeriesAMS
- MQSeriesXRService
- MQSeriesExplorer
- MQSeriesGSKit
- MQSeriesClient
- MQSeriesMan
- MQSeriesMsg
- MQSeriesSamples
- MQSeriesSDK

Procedure

1. Log in as root, or switch to the superuser by using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the server DVD, a network location, or a local file system directory.
3. Run the `mqlicense.sh` script. If you want to view a text-only version of the license, which can be read by a screen reader, type the following message:

```
./mqlicense.sh -text_only
```

The license number is displayed.

You must accept the license agreement before you can proceed with the installation.

4. If this installation is not the first installation of IBM MQ on the system, you must run the **crtmqpkg** command to create a unique set of packages to install on the system. To run the **crtmqpkg** command to run on Linux, you must install the **pax** command and **rpmbuild**, which is located in the `rpm-build` package.

Note: The **crtmqpkg** command is required only if this is not the first installation of IBM MQ on the system. If you have earlier versions of IBM MQ installed on your system, then installing the latest version works correctly if you install it in a different location.

To run the **crtmqpkg** command on a Linux system:

- a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing that uniquely identifies the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

Note: This command creates a full copy of the installation packages in a temporary directory. By default, the temporary directory is located at `/var/tmp`. You must ensure that the system has enough free space before you run this command. To use a different location, you can set the `TMPDIR` environment variable before you run the `crtmqpkg` command. For example:

```
$ TMPDIR=/test ./crtmqpkg suffix
```

- b. Set your current directory to the location specified when the `crtmqpkg` command operation completes successfully. This directory is a subdirectory of the `/var/tmp/mq_rpms` directory, in which the unique set of packages is created. The packages have the `suffix` value contained within the file name. For example, using a suffix of "1":

```
./crtmqpkg 1
```

means there is a subdirectory named `/var/tmp/mq_rpms/1/x86_64`.

The packages are renamed according to the subdirectory, for example:

```
From: MQSeriesRuntime-8.0.0-0.x86_64.rpm  
To: MQSeriesRuntime-1-8.0.0-0.x86_64.rpm
```

5. Install IBM MQ. At a minimum, you must install the `MQSeriesRuntime` and the `MQSeriesServer` components.

- To install to the default location, `/opt/mqm`, use the `rpm -ivh` command to install each component that you require.

For example, to install the runtime and server components to the default location, use the following command:

```
rpm -ivh MQSeriesRuntime-*.rpm MQSeriesServer-*.rpm
```

To install all components to the default location use the following command:

```
rpm -ivh MQSeries*.rpm
```

- To install to a non-default location, use the `rpm --prefix` option. For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must be either an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

For example, enter the following installation path to install the runtime and server components to the `/opt/customLocation` directory on a 64-bit Linux system:

```
rpm --prefix /opt/customLocation -ivh MQSeriesRuntime-*.rpm  
MQSeriesServer-*.rpm
```

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the `setmqenv` or `crtmqenv` command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- If you want to confirm that the installation was successful, you can verify your installation. See [“Verifying a server installation”](#) on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Related information:

setmqinst

Changing the primary installation

Non-interactive installation of IBM MQ server on Linux

You can carry out a non-interactive installation of the IBM MQ server on Linux.

About this task

To carry out a non-interactive installation of IBM MQ, accept the IBM MQ license in non-interactive mode and then follow the interactive installation procedure.

Procedure

1. Log in as root, or switch to the superuser by using the **su** command.
2. You must accept the terms of the license agreement before you can proceed with the installation. To do this run the `mqlicense.sh` script.
The license agreement will be displayed in a language appropriate to your environment and you will be prompted to accept or decline the terms of the license.
If possible, `mqlicense.sh` opens an X-window to display the license.
If you need the license to be presented as text in the current shell, which can be read by a screen reader, type the following command, `mqlicense.sh -text_only`
3. Follow the procedure detailed in “Installing IBM MQ server on Linux” on page 263 or “Installing on Linux Ubuntu or Linux on POWER Systems - Little Endian ” on page 267 as appropriate.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Related information:

setmqinst

Changing the primary installation

Checking the availability of RPM on your Linux Ubuntu machine

You must ensure that RPM is installed on your Linux Ubuntu machine before installing IBM MQ.

About this task

Important: The installation procedure uses the same RPM packages that are used by the other RPM based distributions. Technologies that convert these RPM packages into other forms, such as alien to convert RPMs to Debian packages, are not compatible with the IBM MQ RPM packages and must not be used.

Procedure

1. To determine if the correct RPM package is installed on your system use the following command:

```
dpkg-query -W --showformat '${Status}\n' rpm
```

If you receive a response that is of the form:

```
install ok installed
```

RPM is installed on your system and no further action is required.

If you receive a response that is of the form:

```
unknown ok not-installed
```

RPM is not installed on your system and you must install the RPM package before attempting to install IBM MQ, using the command described in step 2.

2. Run the following command, using root authority. In the example, you obtain root authority using the `sudo` command:

```
sudo apt-get install rpm
```

Attention: If this command does not complete successfully, consult your system administrator for instructions specific to your system on how to install the RPM package.

What to do next

You are now ready to install IBM MQ.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the `rpm` command.

Related information:

setmqinst

Changing the primary installation

Installing on Linux Ubuntu or Linux on POWER Systems - Little Endian

V8.0.0.2

You can install an IBM MQ server on a Linux Ubuntu system, or Linux on POWER Systems - Little Endian system, in accordance with the system requirements web page.

Before you begin

See System Requirements for IBM MQ V8.0 for details of the supported software levels.

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- If you install a copy of IBM MQ server for Linux Ubuntu using Electronic Software Download, obtained from Passport Advantage, you need to uncompress the tar.gz file, and extract the installation files from the tar file, by using the following command:

```
tar -xvf WS_MQ_V8.0_TRIAL_FOR_Linux Ubuntu_ML.tar
```

Important: You must use GNU tar (also known as gtar) to unpack any tar images.

- Ensure that RPM is installed on your system, as RPM is not installed by default on this platform. To determine if the correct RPM package is installed on you system, see “Checking the availability of RPM on your Linux Ubuntu machine” on page 266.
- Once RPM is installed on your system, carry out the following procedure, as root:
 1. Create directory /etc/rpm
 2. Add file /macros, containing the following code, %_dbpath /var/lib/rpm, to the /etc/rpm directory.

Attention: You should only set up a /macros file if you are not already using RPM, as the previous instruction changes the default system-wide RPM database.

About this task

Install the server by using the RPM Package Manager installer to select the components that you want to install. The components and package names are listed in “Choosing what to install” on page 194.

Procedure

1. Open a shell terminal and set your current directory to the location of the installation packages. The location might be the mount point of the server DVD, a network location, or a local file system directory. You must have root authority to run the following commands. You can do so by adding **sudo** before the following commands, or by changing to the root user in the shell with the **su** command.
2. Run the `mqlicense.sh` script. If you want to view a text-only version of the license, which can be read by a screen reader, type the following message:


```
./mqlicense.sh -text_only
```

The license is displayed.

You must accept the license agreement before you can proceed with the installation.

3. If this installation is not the first installation of IBM MQ on the system, you must run the **crtmqpkg** command to create a unique set of packages to install on the system. For the **crtmqpkg** command to run on Linux, you must install the **pax** command and **rpmbuild**, which is located in the rpm package.
 - a. Enter the following command:


```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that uniquely identifies the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

Note: This command creates a full copy of the installation packages in a temporary directory. By default, the temporary directory is located at /var/tmp. You must ensure that the system has enough free space before you run this command. To use a different location, you can set the **TMPDIR** environment variable before you run the **crtmqpkg** command. For example:

```
TMPDIR=/test ./crtmqpkg
```

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a subdirectory of the /var/tmp/mq_rpms directory, in which the unique set of packages

is created. The packages have the *suffix* value contained within the file name. In the following example, the suffix of "1" `./crtmqpkg 1` means that there is a subdirectory named `/var/tmp/mq_rpms/1/i386`.

The packages are renamed according to the subdirectory, for example, on Linux for System x (64-bit):

```
From: MQSeriesRuntime-8.0.0-0.x86_64.rpm
To: MQSeriesRuntime_1-8.0.0-0.x86_64.rpm
```

4. Install IBM MQ. At a minimum, you must install the MQSeriesRuntime and the MQSeriesServer components.

V 8.0.0.2 An additional flag is required when installing on Linux on POWER Systems - Little Endian

- **--ignorearch:** You must include this option to prevent problems with some levels of rpm not recognizing the Linux on POWER Systems - Little Endian architecture

An additional flag is required when installing on Linux Ubuntu:

- **--force-debian:** You must include this option to prevent warning messages from the version of RPM for your platform, which indicates that the RPM packages are not intended to be directly installed using RPM.

To install the IBM MQ Explorer on Linux Ubuntu (x86-64 only):

- a. Install all the components that you want, except for the IBM MQ Explorer component.
- b. Install the IBM MQ Explorer component with the **--nodeps** flag. If you do not include the **--nodeps** flag, the installation fails with a dependency error. The dependency error occurs because the GTK2 packages are not installed by RPM and therefore cannot be found as package dependencies.

If you are installing a subset of components, you must ensure that any dependencies are first installed, as listed in Table 46 on page 270.

Note the following:

- To install to the default location, `/opt/mqm`, use the `rpm -ivh` command to install each component that you require.

To install the runtime and server components to the default location on Ubuntu Linux for System x (64-bit), use the following command:

```
rpm -ivh --force-debian MQSeriesRuntime-*.rpm MQSeriesServer-*.rpm
```

V 8.0.0.2 To install the runtime and server components to the default location on Linux on POWER Systems - Little Endian, use the following command:

```
rpm -ivh --force-debian --ignorearch MQSeriesRuntime-*.rpm MQSeriesServer-*.rpm
```

To install all components to the default location on Linux on POWER Systems - Little Endian use the following command:

```
rpm -ivh --force-debian --ignorearch MQSeries*.rpm
```

- To install to a nondefault location, use the `rpm --prefix` option. For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must be, either an empty directory, the root of an unused file system, or a path that does not exist.

Attention: The length of the path is limited to 256 bytes and must not contain spaces.

V 8.0.0.2 For example, enter the following installation path to install the runtime and server components to the `/opt/customLocation` directory on Linux on POWER Systems - Little Endian :

```
rpm --prefix /opt/customLocation -ivh --force-debian --ignorearch
MQSeriesRuntime-*.rpm MQSeriesServer-*.rpm
```

Table 46 on page 270 lists all available packages on Ubuntu, together with all the associated dependencies.

To install and use the package listed in the *Package Name* column, you must also install the components listed in the *Package Dependencies* column.

Table 46. Package component dependencies

Package Name	Component Function	Dependencies
MQSeriesRuntime	Common function for all other components	None
MQSeriesServer	Queue Manager	MQSeriesRuntime
MQSeriesClient	C IBM MQ client libraries	MQSeriesRuntime
MQSeriesJava	Java and JMS IBM MQ APIs	MQSeriesRuntime
MQSeriesJRE	Java Runtime Environment	MQSeriesRuntime
MQSeriesExplorer	<p>IBM MQ Explorer</p> <p>IBM MQ Explorer is only available on Linux for System x (64-bit).</p> <p>Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2 (or later).</p>	<p>MQSeriesRuntime</p> <p>MQSeriesJRE</p> <p>GTK2 version 2.24-0 or later, including the GTK2 engines that contain the GTK2 themes</p> <p>Bitstream-vera-fonts</p>
MQSeriesGSKit	<p>IBM Global Security Kit</p> <p>Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2 (or later).</p>	<p>MQSeriesRuntime</p> <p>MQSeriesJRE</p>
MQSeriesSDK	Header files and libraries for non-Java APIs	MQSeriesRuntime
MQSeriesMan	UNIX man pages for IBM MQ	MQSeriesRuntime
MQSeriesSamples	IBM MQ application samples	MQSeriesRuntime

Table 46. Package component dependencies (continued)

Package Name	Component Function	Dependencies
MQSeriesMsg_cs MQSeriesMsg_de MQSeriesMsg_es MQSeriesMsg_fr MQSeriesMsg_hu MQSeriesMsg_it MQSeriesMsg_ja MQSeriesMsg_ko MQSeriesMsg_pl MQSeriesMsg_pt MQSeriesMsg_ru MQSeriesMsg_Zh_CN MQSeriesMsg_Zh_TW	Language specific message catalog files	MQSeriesRuntime
MQSeriesFTBase	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesJava MQSeriesJRE
MQSeriesFTLogger	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesServer MQSeriesFTBase MQSeriesJava MQSeriesJRE
MQSeriesFTTools MQSeriesFTAgent	IBM MQ Managed File Transfer components	MQSeriesRuntime MQSeriesFTBase MQSeriesJava MQSeriesJRE
MQSeriesFTService	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesServer MQSeriesFTAgent MQSeriesFTBase MQSeriesJava MQSeriesJRE

Table 46. Package component dependencies (continued)

Package Name	Component Function	Dependencies
MQSeriesAMS	Advanced Message Security component Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2 (or later).	MQSeriesRuntime MQSeriesServer

Results

You have installed the packages you require.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- If you want to confirm that the installation was successful, you can verify your installation. See [“Verifying a server installation”](#) on page 363, for more information.

Related concepts:

[“Multiple installations”](#) on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

[“Choosing a primary installation”](#) on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

[“Uninstalling IBM MQ on Linux”](#) on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Related information:

[setmqinst](#)

[Changing the primary installation](#)

Installing IBM MQ server on Solaris

You can install an IBM MQ server on Solaris either interactively or silently.

Before you begin

- Before you start the installation procedure, make sure that you complete the necessary steps that are outlined in [“Preparing the system”](#) on page 233.

- If you install a copy of IBM MQ server for Solaris using Electronic Software Download, obtained from Passport Advantage, you need to uncompress the tar.gz file, and extract the installation files from the tar file, by using the following command:

```
tar -xvf WS_MQ_V8.0_TRIAL_FOR_SOLARIS_ML.tar
```

Important: You must use GNU tar (also known as gtar) to unpack the tar images.

- If you are using Solaris zones, you have a choice between installing IBM MQ into the global zone, or installing IBM MQ into a non-global zone.

For details on how to install IBM MQ into Solaris zones, see the following technote: WebSphere MQ support position regarding Solaris zones. The technote is applicable to IBM WebSphere MQ Version 7.1 or later with the following changes:

- You do not need the `-G` option on the **pkgadd** command as GSKit is now installed as part of the IBM MQ installation.
- If you install IBM MQ into the global zone for use in sparse zones, you must copy the `/var/mqm` file system into the sparse zone. You must also copy the `/etc/opt/mqm/mqinst.ini` installation entry into the sparse zone.
- Limitations for shared `/usr` file systems: the **dspmqinst** and **dspmqver** commands might report the primary installation incorrectly when compared with the symbolic links in `/usr/bin`. To synchronize the reporting of the primary installation in a Solaris zone and the global zone, run **setmqinst** with the `-i` or `-x` parameter, on the individual zones.
- You cannot change the primary installation within a non-global zone. You must change the primary installation through the global zone, which has the appropriate write access to `/usr/bin`.

About this task

This task describes the installation of the IBM MQ for Solaris server, using the **pkgadd** program. You can choose which components you want to install. The components are listed in “Choosing what to install” on page 194.

Note: If you are installing on the Solaris 11 operating system, ensure that the IPS package (package/svr4) that supports **pkgadd** and equivalent utilities is installed.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the server DVD, a network location, or a local file system directory.
3. Run the `mqlicense.sh` script to accept the license:

```
./mqlicense.sh
```

If you want to view a text-only version of the license, which can be read by a screen reader, type:

```
./mqlicense.sh -text_only
```

The license is displayed. Follow the instructions to accept the license. If you accept the license, the installation continues. If you do not accept the license, you cannot continue the installation process.

4. If this installation is not the first installation on the system, you must run **crtmqpkg** to create a unique set of packages to install on the system:
 - a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing that uniquely identifies the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

The **crtmqpkg** script can use two environment variables which are useful when you are installing from a non-disk media location:

- *CDROOT*, the root of the installation media or downloaded installation files.
- *TMPDIR*, the output location of the modified installation files.

No environment variables are required if you are running the image as `./crtmqpkg`.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a subdirectory of `/var/spool`, in which the unique set of packages is created. The packages have the *suffix* value contained within the file name.
5. Start the installation process:
 - If the installation is the first installation on the system, enter the following command to start the installation process:
`pkgadd -d.`

where `" . "` means use the current directory.
 - If the installation is not the first installation on the system, enter the following command to start the installation process:
`pkgadd mqm-suffix`

where *suffix* is the suffix chosen in the previous step.
6. You are prompted to choose a location for installation.
 - To install to the default location, `/opt/mqm`, enter `y`.
 - To install to a non-default directory, enter `n`. Then, enter the required installation path, and confirm your choice.
7. When the list of components is displayed, enter the numbers of the components that you require, separated by spaces or commas. When you are installing (adding) an IBM MQ component to an existing installation, choose option `yes` when you are asked whether to overwrite.

Note: During the IBM MQ base version installation, you can choose to install all components or a subset of the components. When you install a fix pack, only the currently installed components are upgraded. If, at a later stage, you want to add further IBM MQ components that are not already installed, these components can be installed (added) to the IBM MQ base version only. If your current version of IBM MQ is not the base version, you must first uninstall all the fix packs before you add the required components to the existing installation, and then install the required fix packs. Also, when you are adding IBM MQ components to an existing installation, you must choose option `yes` when you are asked whether to overwrite by the installation process.

8. If the path chosen in step 6 does not exist, you are asked if you want to create it. You must enter `y` to proceed.
9. Answer any questions appropriately for your system. If you are prompted to choose whether to install certain IBM MQ files as `setuid/setgid` files, you must enter `y`.
10. A message is issued when the installation is complete. Enter `q` to exit the **pkgadd** program.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ . For more information, see **setmqenv** and **crtmqenv**.
- If you want to confirm that the installation was successful, you can verify your installation. See “Verifying a server installation” on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Non-interactive installation of IBM MQ server on Solaris”

“Uninstalling IBM MQ on Solaris” on page 394

On Solaris, you can uninstall the IBM MQ server or client using the **pkgrm** command.

Related information:

setmqinst

Changing the primary installation

Non-interactive installation of IBM MQ server on Solaris

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

You can perform a silent installation of IBM MQ. A sample script file called `silent.sh` is supplied in the `silent` directory on the DVD. You can use this script to perform a non-interactive installation that requires no input and shows nothing on the screen. It must be run as root.

The installation script `silent.sh` uses an `admin` file and a response file, both of which are supplied in the `silent` directory. You can use these files as supplied to perform a silent installation of all the components, including all the national language features, to the default location.

Note: If you are installing on the Solaris 11 operating system, ensure that the IPS package (package/svr4) that supports **pkgadd** and equivalent utilities is installed.

Procedure

1. Copy the `silent.sh` script into a writeable directory.
2. If this installation is not the first installation on the system, you must run **crtmqpkg** to create a unique set of packages to install on the system:
 - a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that will uniquely identify the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a sub-directory of `/var/spool`, in which the unique set of packages is created. The packages have the *suffix* value contained within the filename.

Once a new package has been generated for the second installation the `silent.sh` script needs to have its `MQ_PACKAGE_NAME` variable modified so that its value is not `mqm` but the new package name.

Also the `MQ_PACKAGE_LOCATION` variable needs to be modified so that its value is not `$MQ_MEDIA_LOCATION` but the location of the new package (which by default is `/var/spool/pkg`).

3. Optional: If you want to change where the IBM MQ server DVD is mounted, you must update the values in the `silent.sh` script. By default, the script assumes that the server DVD has been mounted at `/CD7FVML`.
4. Optional: If you want to change where the output and logs are written to, update the values in the `silent.sh` script. By default, output and logs are written to the file `/var/tmp/mq.install`.
5. Optional: If you want to install to a non-default location, you must update the `MQ_INSTALLATION_PATH` variable in the `silent.sh` script.

Note:

- The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.
 - If the directory you specified does not exist, the install script creates that directory.
6. Optional: If you want to change the components that are installed, you must edit the response file. A list of all the installable IBM MQ components can be found at: “Choosing what to install” on page 194.

Solaris does not check, during a silent installation, that prerequisite components are installed. You can use the following procedure to create a response file interactively, before using it to install the product. **pkgask** prompts you for the names of the components to install.

- a. Run the **mqlicense.sh** command to accept the license agreement for the product.
- b. **pkgask -d path_to_install_image -r response_file mqm**

The inputs to **pkgask** are the same as those inputs documented for **pkgadd**, but instead of the product being installed a response file is created.

7. Optional: If you have edited the response file, you must then edit the `silent.sh` to use your custom response file.
8. To start the installation, run `silent.sh`.
9. Check the log file for any errors.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- If you want to confirm that the installation was successful, you can verify your installation. See [“Verifying a server installation”](#) on page 363, for more information.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Installing IBM MQ server on Solaris” on page 272

You can install an IBM MQ server on Solaris either interactively or silently.

“Uninstalling IBM MQ on Solaris” on page 394

On Solaris, you can uninstall the IBM MQ server or client using the **pkgrm** command.

Related information:

setmqinst

Changing the primary installation

Installing IBM MQ server on Windows

This topic describes how to install IBM MQ server on Windows systems by using the Launchpad. This procedure can be used for installing a first or a subsequent installation.

Installing using the Launchpad

About this task

These instructions cover how to display the installation Launchpad window. You can use the launchpad to make a compact, typical, or custom installation of IBM MQ. You can reuse the launchpad multiple times to install further installations. It automatically selects the next available installation name, instance, and location to use. To view all the installation types and the features that are installed with each option, see “Planning your installation on Windows systems” on page 215.

Note that if you have previously uninstalled IBM MQ from your system (see “Uninstalling IBM MQ on Windows systems” on page 395), some configuration information might remain, and some default values might be changed.

Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate **setup.exe** in the base directory of the IBM MQ installation image.
 - From a DVD, this location might be `E:\setup.exe`
 - From a network location, this location might be `m:\instmq\setup.exe`
 - From a local file system directory, this location might be `C:\instmq\setup.exe`
3. Double-click the **Setup** icon to start the installation process. It is possible to run either by:
 - Running `setup.exe` from the command prompt. Or
 - Double-clicking `setup.exe` from Windows Explorer.

If you are installing on a Windows system with UAC enabled, accept the Windows prompt to allow the launchpad to run as elevated. During installation, you might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the installation to continue.

The IBM MQ Installation Launchpad window is displayed.

4. Continue to follow the Launchpad instructions as shown on screen.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH\bin\setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see [“Verifying a client installation”](#) on page 374.

Related concepts:

[“Modifying your installation”](#) on page 293

Find out how to modify an IBM MQ server installation interactively using the launchpad or non-interactively using **msiexec**.

[“Post installation tasks”](#) on page 294

Find out the tasks that can be undertaken when IBM MQ has been successfully installed. Begin by following the related pages:

Related tasks:

[“Advanced installation using msiexec”](#) on page 279

[“Uninstalling IBM MQ on Windows systems”](#) on page 395

You can uninstall the IBM MQ MQI clients and servers on Windows systems by using the control panel, the command line (**msiexec**), **MQParms**, or by using the installation media, in which case you can optionally remove queue managers as well.

Installing IBM MQ using SCCM

You can install IBM MQ using the System Center Configuration Manager (SCCM).

Procedure

1. Create a folder on the SCCM server and copy the IBM MQ DVD image into that folder.
2. Make the folder the X drive of the SCCM client systems.
3. Using the Configuration Manager Console for SCCM, create a package:
 - a. Click **Computer Management > Software Distribution > Packages**.
 - b. Right-click on **Packages** and click **New > Package from Definition**.
 - c. In the wizard, select the WebSphere MQ.SMS file from the folder that you copied the IBM MQ DVD image into. If you chose to use a drive letter other than X for the client systems, you must modify the commands in the WebSphere MQ.SMS file to use the appropriate drive letter.
4. Optional: To install IBM MQ to a non-default location, you must add the following two parameters to the command line under Program Properties:

```
PGMFOLDER=" MQ_INSTALLATION_PATH "  
DATFOLDER=" MQ_INSTALLATION_PATH\data"
```

where *MQ_INSTALLATION_PATH* refers to the path where you want to install IBM MQ.

5. Under the package you created, Right click **Distribution Points** and then click **New distribution Points**.
6. In the wizard, select the SCCM server as the distribution point.
7. Using the Configuration Manager Console for SCCM, distribute the software:
 - a. Right-click the package you created and click **Distribute > Software**.
 - b. In the wizard, select **All Systems for Collections**.

8. Once the Advertisement Status for the package shows as Received, manually initiate the IBM MQ unattended installation:
 - a. Log on to the SCCM client system
 - b. Click **Control Panel > Run Advertised Program** and select the package to run

Advanced installation using msixec

Before you begin

If you are running IBM MQ on Windows systems with User Account Control (UAC) enabled, you must invoke the installation with elevated privileges. If you are using the Command prompt or MQ Explorer elevate privileges by using a right-click to start the program and selecting Run as administrator. If you try to run **msixec** without using elevated privileges, the installation fails with an error of AMQ4353 in the installation log.

About this task

IBM MQ on Windows uses the MSI technology to install software. MSI provides both an interactive installation and a non interactive installation. An interactive installation displays panels and ask questions.

The **msixec** command uses parameters to give MSI some or all of the information that can also be specified through panels during an interactive installation. This means that a user can create a reusable automated or semi-automated installation configuration. Parameters can be given through the command line, a transform file, a response file, or a combination of the three.

Procedure

To install using **msixec**, at the command line, enter the **msixec** command in the following format:
`msixec parameters[USEINI="response-file"] [TRANSFORMS="transform_file"]`

Where:

parameters

are either command-line parameters preceded by a / character, or property=value pairs (if using both forms of parameter always put the command-line parameters first). For further information, see "Specifying command line parameters with **msixec**" on page 280, which contains a link to the web site that lists all the command line parameters that are available.

For an unattended installation, you must include the /q or /qn parameter in the command line. Without this parameter, the installation is interactive.

Note: You must include the /i parameter and the file location of the IBM MQ installer package.

response-file

is the full path and file name of the file that contains the [Response] stanza and the required property=value pairs, for example C:\MyResponseFile.ini. An example response file, Response.ini, is supplied with IBM MQ. This file contains default installation parameters. For further information, see "Using a response file with **msixec**" on page 281.

transform_file

is the full path and file name of a transform file. For further information, see "Using transforms with **msixec**" on page 286 and "Multiple installation using MSI Instance ID" on page 285.

Note: For a silent installation to succeed, the AGREETOLICENSE=?YES? property must be defined either on the command line or in the response file.

Results

After the command has been entered, the command prompt immediately reappears. IBM MQ is installing as a background process. If you have entered parameters to produce a log, check this file to see how the installation is progressing. If the installation completes successfully, you see the message `Installation operation completed successfully` in the log file.

Specifying command line parameters with `msiexec`: About this task

The `msiexec` command can accept two types of parameters on the command line, as follows:

- Standard command line parameters, preceded by a `/` character.
For a table of the `msiexec` command line parameters, see the MSDN Command-Line Options web page.
- Property=value pair parameters on the command line. All the parameters available for use in a response file can be used on the command line, for a list of these, see Table 48 on page 282. In addition there are some extra property=value pair parameters that are only for use on the command line, for details see Table 47 on page 281.

When using the property=value pair parameters note that:

- Property strings must be in uppercase.
- Value strings are not case-sensitive, except for feature names. You can enclose value strings in double quotation marks. If a value string includes a blank, enclose the blank value string in double quotation marks.
- For a property that can take more than one value, use the format:
`ADDLOCAL="Server,Client"`
- For properties taking paths and filenames, for example `PGMFOLDER`, you must supply the paths as absolute paths and not relative; that is, `C:\folder\file` and not `.\folder\file`.

When using property=value pair and command line parameters with the `msiexec` command, enter command line parameters first.

If a parameter is specified both on the command line and in a response file, the setting on the command line takes precedence.

Example

Here is an example of a typical `msiexec` command. All parameters, separated by one or more spaces, must be typed on the same line as the `msiexec` call.

```
msiexec
/i "<path>\MSI\IBM WebSphere MQ.msi"
/l*v c:\install.log
/q
TRANSFORMS="1033.mst"
AGREETOLICENSE="yes"
ADDLOCAL="Server"
```

Here is an example of a typical `msiexec` command when you are installing a second copy of IBM WebSphere MQ Version 7.5, or later. All parameters, separated by one or more spaces, must be typed on the same line as the `msiexec` call.

```
msiexec
/i "<path>\MSI\IBM WebSphere MQ.msi"
/l*v c:\install.log
/q
TRANSFORMS=":InstanceId2.mst;1033.mst"
AGREETOLICENSE="yes"
ADDLOCAL="Server"
MSINewInstance=1
```


Where /! *v c:\install.log writes installation log to file c:\install.log.

The following table shows the parameters which can only be provided on the command line and not in a response file.

Table 47. *msiexec property=value parameters*

Property	Values	Meaning
USEINI	<i>path \ file_name</i>	Use the specified response file. See “Using a response file with msiexec”
SAVEINI	<i>path \ file_name</i>	Generate a response file during installation. The file contains those parameters selected for this installation that a user might make during an interactive installation.
ONLYINI	1 yes ""	1, yes or any value other than null. End the installation before updating the target system, but after generating a response file, if this is specified. "". Continue the installation and update the target system (the default).
TRANSFORMS	:InstanceId <i>x.mst</i> <i>path \ file_name</i> :InstanceId <i>x.mst</i> ; <i>path \ file_name</i>	The :InstanceId <i>x.mst</i> value is only required for a subsequent installation of IBM MQ. The <i>path \ file_name</i> specifies what transform (.mst) files must be applied to the product. For example, "1033.mst" specifies the supplied U.S. English transform file.
MSINewInstance	1	This property is only required for subsequent installations of IBM MQ

Using a response file with msiexec: About this task

You can use the **msiexec** command with a parameter which specifies additional properties that are defined in a response file. You can combine the msiexec command-line parameters described in “Specifying command line parameters with msiexec” on page 280.

A response file is an ASCII text file, with a format like a Windows .ini file, that contains the stanza [Response]. The [Response] stanza contains some or all the parameters that would normally be specified as part of an interactive installation. The parameters are given in a property=value pair format. Any other stanzas in the response file are ignored by **msiexec**. An example response file, Response.ini, is supplied with IBM MQ. It contains the default installation parameters.

Procedure

A typical example of an msiexec command is: `msiexec /i "path\MSI\IBM WebSphere MQ.msi" /! *v c:\install.log TRANSFORMS= "1033.mst" USEINI= "C:\MQ\Responsefile"`

If a parameter is specified both on the command line and in a response file, the setting on the command line takes precedence. All the parameters available for use in a response file can also be used on the command line, for a list of these see Table 48 on page 282.

In the response file, all text is in English, and comments begin with a ; character.

For information about creating a response file, see “Creating a response file” on page 287.

Example

An example of a typical response file:

```

[Response]
PGMFOLDER="c:\mqm"
DATFOLDER="c:\mqm\data"
LOGFOLDER="c:\mqm\log"
AGREETOLICENSE="yes"
LAUNCHWIZ=""
WIZPARMFILE="d:\MQParms.ini"
ADDLOCAL="Server,Client"
REMOVE="Toolkit"

```

Table 48. Response file parameters

Property	Values	Meaning
PGMFOLDER	<i>path</i>	Folder for the IBM MQ program files. For example, c:\mqm.
DATFOLDER	<i>path</i>	Folder for the IBM MQ data files. For example, c:\mqm\data. Note: Multiple installations of IBM MQ all use the same DATFOLDER .
LOGFOLDER	<i>path</i>	Folder for the IBM MQ queue manager log files. For example, c:\mqm\log. Note: Multiple installations of IBM MQ all use the same LOGFOLDER .
USERCHOICE	0 no	If the command line or response file specifies parameters to install features, a dialog can be displayed to prompt the user to accept the preselected options, or review and possibly change them. 0 or no. Suppresses display of the dialog. Anything else. Dialog is displayed. Not used for a silent installation.
AGREETOLICENSE	yes	Accept the terms of the license. Set to yes before a silent installation. If the installation is not silent, this parameter is ignored.
KEEPQMDATA	keep delete	If the Server feature is to be uninstalled, whether to delete any existing queue managers. delete removes any existing queue managers. keep, or any other value, keeps them. Note: This property is only valid on a final server uninstall. If used on any other server uninstalls this property is ignored.

Table 48. Response file parameters (continued)

Property	Values	Meaning
LAUNCHWIZ	0 1 yes no ""	<p>0 or no. Do not launch the Prepare IBM MQ wizard after IBM MQ is installed.</p> <p>1 or yes. Launch the Prepare IBM MQ wizard if the Server feature is installed.</p> <p>"". Launch the Prepare IBM MQ wizard to install the Server (the default).</p> <p>If this option is to launch the Prepare IBM MQ wizard, you can specify the WIZPARMFILE, either in this file, or on the command line.</p> <p>The Prepare IBM MQ wizard must be run to make your IBM MQ installation operational. If you choose not to launch it here, you must run it before using IBM MQ.</p>
WIZPARMFILE	<i>path \ file_name</i>	When specified, the file that contains the parameters to pass to the Prepare IBM MQ wizard when it is launched. These are in the [Services].
ADDLOCAL	<i>feature, feature, All ""</i>	<p>A comma-separated list of features to install locally. For a list of valid feature names, see "IBM MQ features for Windows systems" on page 207.</p> <p>All installs all features</p> <p>"" installs the typical features. If you do not want a feature use REMOVE=" <i>feature</i> "</p> <p>Note: If this is a new installation the typical features (Server, Explorer, Java Messaging and SOAP Transport, and Development Toolkit) are installed by default irrespective of the feature list provided in the ADDLOCAL property. If you do not want a feature use REMOVE="feature"</p>
REMOVE	<i>feature, feature, All ""</i>	<p>A comma-separated list of features to remove. For a list of valid feature names, see "IBM MQ features for Windows systems" on page 207.</p> <p>All uninstalls all features</p> <p>"" uninstalls no features (the default).</p>

Table 48. Response file parameters (continued)

Property	Values	Meaning
STARTSERVICE	0 no ""	<p>0 or no. Do not start the IBM MQ Service at the end of installation.</p> <p>"" (the default). Start the IBM MQ Service at the end of installation if it was running at the start, or if this is a new installation.</p> <p>Anything else. Start the Service at the end of the installation.</p> <p>Ignored if the server feature is not installed.</p> <p>If you do not start the IBM MQ Service, IBM MQ will not be operational and queue managers will not start. You must run the Prepare IBM MQ wizard for the service to be correctly configured.</p> <p>This parameter is only valid if LAUNCHWIZ is set to no.</p>
STARTTASKBAR	0 no ""	<p>0 or no. Do not start the IBM MQ taskbar application at the end of installation.</p> <p>"" (the default). Start the IBM MQ taskbar application at the end of installation if it was running at the start, or if this is a new installation.</p> <p>Anything else. Start the taskbar application at the end of the installation.</p> <p>Ignored if the server feature is not installed.</p> <p>This parameter is only valid if LAUNCHWIZ is set to no.</p>
INSTALLATIONDESC	?Description of installation?	Sets the installation description from the command line. Subject to the documented installation description length limitations
INSTALLATIONNAME	[INSTALLATION0,]?Name?	<p>Sets the installation name from the command line. Subject to the documented installation name character and length limitations.</p> <p>Note: Supply INSTALLATION0,Name only when upgrading from versions of IBM MQ before Version 7.5.</p>
MAKEPRIMARY	0 1 ""	<p>Makes the installation primary, if possible, or removes the primary flag. 1 = Make primary, 0 = Make non-primary, - use default algorithm</p> <p>Note: This option is ignored if a version of IBM MQ before Version 7.5 is installed, or if another installation of IBM WebSphere MQ Version 7.5, or later, is present and set as the primary.</p>

Related tasks:

“Multiple installation using MSI Instance ID”

This topic describes how to choose the MSI instance ID you require for non-interactive multiple installations.

“Creating a response file” on page 287

A response file is used with **msiexec**. You can create it in three ways.

“Using the MQParms command” on page 287

Related reference:

“Using transforms with msiexec” on page 286

Multiple installation using MSI Instance ID:

This topic describes how to choose the MSI instance ID you require for non-interactive multiple installations.

About this task

In order to support non-interactive multiple installations, you need to find out whether the instance ID you want to use is already in use or not and choose the appropriate one. For each installation media (for example, each client and server), Instance ID 1 is the default ID which is used for single installations. If you want to install alongside Instance ID 1 you need to specify which instance you want to use. If you have already installed instance 1, 2, and 3 then you need to find out what the next available instance is, for instance, Instance ID 4. Similarly, if instance 2 has been removed, you need to find out that there is a gap that can be reused. You can find out which Instance ID is currently in use by using the **dspmqinst** command.

Procedure

1. Type **dspmqinst** to find a free MSI Instance in the media being installed by reviewing the MSIMedia and MSIInstanceId values for the versions already installed. For example:

```
InstName: Installation1
InstDesc:
Identifier: 1
InstPath: C:\Program Files\IBM\WebSphere MQ
Version: 8.0.0.0
Primary: Yes
State: Available
MSIProdCode: {74F6B169-7CE6-4EFB-8A03-2AA7B2DBB57C}
MSIMedia: 8.0 Server
MSIInstanceId: 1
```

2. If MSI Instance ID 1 is in use and you want to use MSI Instance ID 2, the following parameters must be added to the msiexec call:

```
MSINEWINSTANCE=1 TRANSFORMS=:InstanceId2.mst
```

What to do next

For multiple installations, the **INSTALLATIONNAME** or **PGMFOLDER** must be supplied as an additional parameter on any non-interactive installation command. Supplying the **INSTALLATIONNAME** or **PGMFOLDER** ensures that you do not work with the wrong installation in case you omit or incorrectly specify the **TRANSFORMS** parameter.

Using transforms with msixec:

MSI can use transforms to modify an installation. During IBM MQ installation, transforms can be used to support different national languages. IBM MQ is supplied with transform files in the \MSI folder of the Server image. These files are also embedded in the IBM MQ Windows installer package, IBM WebSphere MQ.msi.

On the **msiexec** command line, you can specify the required language by using the TRANSFORMS property in a property=value pair. For example:

```
TRANSFORMS="1033.mst"
```

You can also specify the full path and file name of the transform file. Again, the quotation marks surrounding the value are optional. For example:

```
TRANSFORMS="D:\Msi\1033.mst"
```

Table 49 shows the locale identifier, language, and the transform file name to use in the **msiexec** command line.

You might need to merge transforms to install multiple installations of the same version, for example:

```
TRANSFORMS=":InstanceId2.mst;D:\Msi\1033.mst"
```

You can also specify the required language by using the MQLANGUAGE property with the **MQParms** command. For information about the msiexec property=value parameters, see "MQParms parameter file" on page 289.

Parameters

Table 49. Supplied transform files for various language support. This table shows the supplied transform files, the resulting language, and the numeric value to use in the **msiexec** command line.

Language	Transform File name	Value
U.S. English	1033.mst	1033
German	1031.mst	1031
French	1036.mst	1036
Spanish	1034.mst	1034
Italian	1040.mst	1040
Brazilian Portuguese	1046.mst	1046
Japanese	1041.mst	1041
Korean	1042.mst	1042
Simplified Chinese	2052.mst	2052
Traditional Chinese	1028.mst	1028
Czech	1029.mst	1029
Russian	1049.mst	1049
Hungarian	1038.mst	1038
Polish	1045.mst	1045

Creating a response file:

A response file is used with **msiexec**. You can create it in three ways.

About this task

A response file is used with the **msiexec** command, for further information see “Using a response file with msiexec” on page 281.

Procedure

There are three ways to create a response file for installation:

- Copy and edit the file `Response.ini` that is supplied on the IBM MQ Windows Server CD, using an ASCII file editor.
- Create your own response file using an ASCII file editor.
- Use the **msiexec** command with the **SAVEINI** (and optionally, the **ONLYINI**) command line parameters to generate a response file that contains the same installation options. See Table 47 on page 281.

Example

A typical example of using **msiexec** with the **SAVEINI** parameter is here:

```
msiexec /i "path\IBM WebSphere MQ.msi" /q SAVEINI="response_file"  
TRANSFORMS="1033.mst" AGREETOLICENSE="yes"
```

Using the MQParms command:

Before you begin

You can use the **MQParms** command to invoke installation or uninstallation. This command can use parameters on a command line, or those specified in a parameter file. The parameter file is an ASCII text file that contains the parameter values that you want to set for the installation. The **MQParms** command takes the specified parameters and generates the corresponding **msiexec** command line.

This means that you can save all the parameters that you want to use with the **msiexec** command in a single file.

If you are running IBM MQ on Windows systems with User Account Control (UAC) enabled, you must invoke the installation with elevated privileges. If you are using the Command prompt or MQ Explorer elevate privileges by using a right-click to start the program and selecting **Run as administrator**. If you try to run the **MQParms** program without using elevated privileges, the installation fails with an error of AMQ4353 in the installation log.

For silent operations, this must include the **/q** or **/qn** parameter, either on the command line, or in the [MSI] stanza of the parameter file. You must also set the **AGREETOLICENSE** parameter to "yes".

You can specify many more parameters in the parameter file that you use with the **MQParms** command than you can in the response file that you use directly with the **msiexec** command. Also, as well as parameters that the IBM MQ installation uses, you can specify parameters that can be used by the Prepare IBM MQ wizard.

If you do not complete the Prepare WebSphere MQ Wizard directly after IBM MQ installations or if for any reason your machine is rebooted between completing IBM MQ installation and completing the Prepare WebSphere MQ Wizard, ensure that the wizard is run with Administrator privilege afterward, otherwise the installation is incomplete, and might fail. You might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the wizard to continue

An example of the file MQParms.ini is supplied with IBM MQ. This file contains default installation parameters.

There are two ways to create a parameter file for installation:

- Copy and edit the file MQParms.ini that is supplied with the product, using an ASCII file editor.
- Create your own parameter file using an ASCII file editor.

About this task

To invoke installation using the MQParms command:

Procedure

1. From a command line, change to the root folder of the IBM MQ Server DVD (that is, the location of the file MQParms.exe).
2. Enter the following command:

```
MQParms parameter_file parameters ]
```

where:

parameter_file

is the file that contains the required parameter values. If this file is not in the same folder as MQParms.exe, specify the full path and file name. If you do not specify a parameter file, the default is MQParms.ini. For a silent install, the MQParms_silent.ini parameter file can be used. For further details, see “MQParms parameter file” on page 289.

parameters

are one or more command-line parameters, for a list of these, see the MSDN Command-Line Options web page.

Example

A typical example of an MQParms command is:

```
MQParms "c:\MyParamsFile.ini" /!*v c:\install.log
```

A typical example of an MQParms command when you are installing a second copy of IBM MQ Version 7.5 is:

```
MQParms "c:\MyParamsFile.ini" /!*v c:\install.log TRANSFORMS=":InstanceId2.mst;1033.mst" MSINEWINSTANCE=1
```

Alternatively, TRANSFORMS and MSINEWINSTANCE can be specified in the MSI stanza of the parameter file.

If you specify a parameter both on the command line and in the parameter file, the setting on the command line takes precedence.

If you specify a parameter file, you might want to run the encryption utility before you use the MQParms command (see “Encrypting a parameter file” on page 292).

If you do not specify /i, /x, /a, or /j, MQParms defaults to standard installation using the IBM MQ Windows Installer package, IBM IBM WebSphere MQ.msi. That is, it generates the following part of the command line:

```
/i " current_folder \MSI\IBM WebSphere MQ.msi"
```

If you do not specify a WIZPARMFILE parameter, MQParms defaults to the current parameter file. That is, it generates the following part of the command:

```
WIZPARMFILE=" current_folder \ current_parameter_file "
```


MQParms parameter file:

A parameter file is an ASCII text file that contains sections (stanzas) with parameters that can be used by the MQParms command. Typically, this is an initialization file such as MQParms.ini.

The MQParms command takes parameters from the following stanzas in the file:

[MSI] Contains general properties related to how the MQParms command runs and to the installation of IBM MQ.

The properties that you can set in this stanza are listed in “Advanced installation using msixec” on page 279, and Table 50.

[Services]

Contains properties related to IBM MQ account configuration, in particular, the user account required for IBM MQ Services. If you are installing IBM MQ on a network where the domain controller is on a Windows 2003 server, you probably need details of a special domain account. For further information, see “Configure IBM MQ accounts” on page 298 and “Configuring IBM MQ with the Prepare IBM MQ wizard” on page 295.

The properties that you can set in this stanza are listed in Table 52 on page 291.

MQParms ignores any other stanzas in the file.

The stanza parameters are in the form `property=value`, where `property` is always interpreted as uppercase, but `value` is case sensitive. If a value string includes a blank, it must be enclosed in double quotation marks. Most other values can be enclosed in double quotation marks. Some properties can take more than one value, for example:

```
ADDLOCAL="Server,Client"
```

To clear a property, set its value to an empty string, for example:

```
REINSTALL=""
```

The following tables show the properties that you can set. The default is shown in bold.

For the [MSI] stanza, you can enter standard MSI command line options and properties. For example:

- /q
- ADDLOCAL="server"
- REBOOT=Suppress

Refer to Table 50, Table 51 on page 290, and Table 52 on page 291 for the properties used to install IBM MQ.

Table 50 shows additional properties in the stanza that affect how the MQParms command runs, but that do not affect the installation.

Table 50. Properties used by MQParms in the MSI stanza

Property	Values	Description
MQPLOG	<i>path file_name</i>	MQParms generates a text log file with the specified name and location.

Table 50. Properties used by MQParms in the MSI stanza (continued)

Property	Values	Description
MQPLANGUAGE	system user <i>transform_value</i> existing	The installation language. system. Install using the language of the default system locale (the default). user. Install using the language of the default locale of the user. <i>transform_value</i> . Install using the language specified by this value. See Table 51. existing. If MQ already exists on the system, the same language will be used by default, otherwise system is used.
MQPSMS	0 no	0 or no. MQParms does not wait for the <code>msiexec</code> command to end (the default). Any other value. MQParms waits for the <code>msiexec</code> command to end.
MQPINUSE	0 1	If MQPINUSE is set to 1, MQParams continues installing even if IBM MQ files are in use. If this option is used a reboot will be required to complete the installation.

Table 51. Valid values for the MQPLANGUAGE property

Language	Valid values		
U.S. English	English	en_us	1033
German	German	de_de	1031
French	French	fr_fr	1036
Spanish	Spanish	es_es	1034
Italian	Italian	it_it	1040
Brazilian Portuguese	Brazilian Portuguese	pt_br	1046
Japanese	Japanese	ja_jp	1041
Korean	Korean	ko_kr	1042
Simplified Chinese	Simplified Chinese	zh_cn	2052
Traditional Chinese	Traditional Chinese	zh_tw	1028
Czech	Czech	cs_cz	1029
Russian	Russian	ru_ru	1049
Hungarian	Hungarian	hu_hu	1038
Polish	Polish	pl_pl	1045

For the [Services] stanza, you can enter parameters in property=value format. You might want to encrypt the values in this stanza. See “Encrypting a parameter file” on page 292.

Table 52. Properties used in the Services stanza

Property	Values	Description
USERTYPE	local domain onlydomain	<p>The type of user account to use:</p> <p>local Creates a local user account.</p> <p>domain Creates a local user account. If this does not have the required security authorities, it uses the domain user account specified by DOMAINNAME, USERNAME, and PASSWORD.</p> <p>onlydomain Does not create a local user account, but immediately uses the domain user account specified by DOMAINNAME, USERNAME and PASSWORD. If any of these three properties are missing, a USERTYPE of local is assumed.</p> <p>The properties DOMAINNAME, USERNAME, and PASSWORD are required if USERTYPE is set to onlydomain.</p>
DOMAINNAME	<i>domain_name</i> ¹	<p>The domain for the domain user account.</p> <p>Required if USERTYPE is set to domain or onlydomain.</p>
USERNAME	<i>user_name</i> ¹	<p>The user name for the domain user account.</p> <p>Required if USERTYPE is set to domain or onlydomain..</p>
PASSWORD	<i>password</i> ¹	<p>The password for the domain user account.</p> <p>Required if USERTYPE is set to domain or onlydomain.</p>
<p>1. Do not enclose this value in double quotation marks.</p>		

A typical example of a parameter file is:

```
[MSI]
MQPLANGUAGE=1033
MQPLOG=%temp%\MQParms.log
MQPSMS=no
ADDLOCAL=Server
/m miffile
REMOVE=""
/l*v c:\install.log

[Services]
USERTYPE=domain
DOMAINNAME=mqm*df349edfcab12
USERNAME=mqm*a087ed4b9e9c
PASSWORD=mqm*d7eba3463bd0a3
```

Encrypting a parameter file:

About this task

Use the `setmqpw` utility to encrypt the `DOMAINNAME`, `USERNAME`, and `PASSWORD` values in the `[Services]` stanza of a parameter file, if they are not already encrypted. (These values might be encrypted if you have run the utility before.) `setmqpw` will also encrypt the `QMGRPASSWORD` and `CLIENTPASSWORD` values in the `[SSLMigration]` stanza of a parameter file.

This encryption means that, if you need a special domain account to configure IBM MQ (see “Configure IBM MQ accounts” on page 298), or you need to keep key database passwords secret, details are kept secure. Otherwise, these values, including the domain account password, flow across the network as clear text. You do not have to use this utility, but it is useful if security in your network is an issue.

To run the script:

Procedure

1. From a command line, change to the folder that contains your parameter file.
2. Enter the following command:

```
CD_drive:\setmqpw
```

Note: You can run the command from a different folder, by entering the following command, where *parameter_file* is the full path and file name of the parameter file:

```
CD_drive:\setmqpw parameter_file
```

Results

If you view the resulting parameter file, the encrypted values start with the string `mqm*`. Do not use this prefix for any other values; passwords or names that begin with this prefix are not supported.

The utility creates a log file, `setmqpw.log`, in the current directory. This file contains messages related to the encryption process. When encryption is successful, messages are similar to:

```
Encryption complete  
Configuration file closed  
Processing complete
```

What to do next

After you encrypt the parameter file, you can use it in the normal way with the `MQParms` command (see “Using the `MQParms` command” on page 287).

Modifying your installation

Find out how to modify an IBM MQ server installation interactively using the launchpad or non-interactively using msixec.

Related concepts:

“Silently modifying an IBM MQ server installation using msixec” on page 294

Related tasks:

“Modifying the installation using IBM MQ Installation Launchpad”

Modifying the installation using IBM MQ Installation Launchpad: Before you begin

To modify an installation, some features of IBM MQ must already be installed.

About this task

To remove or install IBM MQ features follow the instructions. This procedure is the only way to interactively remove or install IBM MQ features on Windows Server 2008:

Procedure

1. Insert the IBM MQ for Windows Server DVD into the DVD drive.
2. If autorun is installed, the installation process starts.
Otherwise, double-click the **Setup** icon in the root folder of the DVD to start the installation process. The WebSphere MQ Installation Launchpad window is displayed.
3. Click the **IBM MQ Installation** option.
4. Click **Launch IBM MQ Installer**. Wait until the IBM MQ Setup window is displayed with a welcome message.
5. If you have multiple installations on your system, you must choose the installation you want to modify. Do this by selecting the **Maintain or upgrade an existing instance** option and choosing the appropriate instance. If you are upgrading a IBM WebSphere MQ Version 7.0.1 installation (or earlier) to Version 7.1.0, and you already have a Version 7.1.0 or greater installation, you need to select **Install a new instance**. A subsequent panel then allows you to choose the installation you would like to upgrade.
6. Click **Next** to continue. The Program Maintenance panel is displayed.
7. Select **Modify**, then click **Next**.
The Features panel is displayed.
8. Click the + symbol next to a feature to show any dependent features (subfeatures).
9. To change the installation of a feature:
 - a. Click the symbol next to the feature name to display a menu.
 - b. Select the required option from:
 - Install this feature
 - Install this feature and all its subfeatures (if any)
 - Do not install this feature (remove if already installed)

The symbol next to the feature name changes to show the current installation option.
10. To remove a feature stop the IBM MQ instance If you do not do this, you receive an error message.
11. When your selections are complete, click **Next**. IBM MQ installation begins.

What to do next

After modifying the installation, you might need to run **setmqenv** again as described in *What to do next* in “Installing IBM MQ server on Windows” on page 277.

Silently modifying an IBM MQ server installation using msiexec:

To silently modify an installation using msiexec, set the ADDLOCAL parameter to include the features you want to add, and set the REMOVE parameter to the features you want to remove.

For example if you use ADDLOCAL="JavaMsg" and REMOVE="" it modifies the installation to include the Java Messaging and Web Services feature.

```
msiexec /i {PRODUCT CODE} /q ADDLOCAL="JavaMsg" REMOVE="" INSTALLATIONNAME="Installation1"
```

The instructions for msiexec begin here: “Advanced installation using msiexec” on page 279

Post installation tasks

Find out the tasks that can be undertaken when IBM MQ has been successfully installed. Begin by following the related pages:

Related concepts:

“Configuring an IBM MQ Server”

“Configure IBM MQ accounts” on page 298

The IBM MQ service and queue managers check that any users attempting to access queue managers or queue manager resources such as queues, have the permission to access them.

“Using IBM MQ Remotely” on page 295

“Using the Help Center” on page 303

Related tasks:

“Configuring IBM MQ with the Prepare IBM MQ wizard” on page 295

“Using the Default Configuration wizard” on page 302

“Checking for problems after installing” on page 297

These are optional tasks that you can use to check the installation if you believe there was a problem, or to verify installation messages after an unattended (silent) installation for example.

Related reference:

“Using the Welcome to IBM MQ Explorer Content view page” on page 303

The Welcome to IBM MQ Explorer Content view page points you to any relevant applications, documentation, tutorials, and education. This page is displayed the first time you launch IBM MQ Explorer.

Configuring an IBM MQ Server: After installing IBM MQ, it is necessary to configure it. The configuration described in this section is for an environment that uses TCP/IP. The configuration procedure is the same for environments that use other communications protocols (for example, SNA, SPX, or NetBIOS). However, not all of the functions and facilities of IBM MQ for Windows are available in these environments. The items that are **not** available are:

- IBM MQ Postcard
- IBM MQ Explorer

If you are setting up IBM MQ for use with the Microsoft Cluster Service (MSCS), see Supporting the Microsoft Cluster Service (MSCS) for more information.

Using IBM MQ Remotely: If you are connecting to a Windows machine using either Terminal Services or a Remote Desktop Connection and you have problems creating, starting or deleting a queue manager this might be because of the user access **Create global objects**.

The **Create global objects** user access limits the users authorized to create objects in the global namespace. In order for an application to create a global object, it must either be running in the global namespace, or the user under which the application is running must have the **Create global objects** user access applied to it.

When you connect remotely to a Windows machine using either Terminal Services or Remote Desktop Connection, applications run in their own local namespace. If you attempt to create or delete a queue manager using IBM MQ Explorer or the **crtmqm** or **dltmqm** command, or to start a queue manager using the **strmqm** command, it results in an authorization failure. This creates an IBM MQ FDC with Probe ID XY132002.

Starting a queue manager using the IBM MQ Explorer, or using the **amqmdain qmgr start** command works correctly because these commands do not directly start the queue manager. Instead the commands send the request to start the queue manager to a separate process running in the global namespace.

If you need to perform any of these operations on a queue manager when connected remotely to a Windows machine, you must have the **Create global objects** user access. For information on how to assign a user this access, see your operating system documentation.

Administrators have the **Create global objects** user access by default, so if you are an administrator you can create and start queue managers when connected remotely without altering your user rights.

Configuring IBM MQ with the Prepare IBM MQ wizard:

About this task

The Prepare IBM MQ wizard helps you to configure IBM MQ files and a user account for your network, and migrate any queue managers and data from a previous installation. You must run the wizard to configure the IBM MQ Service before you can start any queue managers.

The Prepare IBM MQ wizard window is displayed when IBM MQ installation completes. Follow the instructions given by the wizard to configure IBM MQ. At any time while the wizard is running you can click **More Information** in the wizard to view online help about the task you are doing.

On Windows systems, you must do this task under a Windows administrator account, or domain administrator account in case your workstation is a member of a Windows domain.

On Windows systems with UAC enabled, if you do not complete the Prepare WebSphere MQ Wizard directly after WebSphere MQ is installed, or if for any reason your machine is rebooted between completing IBM MQ installation and completing the Prepare WebSphere MQ Wizard, you must accept the Windows prompt when it appears to allow the wizard to run as elevated.

Procedure

1. When the IBM MQ installation completes, the Prepare WebSphere MQ Wizard window is displayed with a welcome message. To continue, click **Next**
2. If you have run the Prepare IBM MQ wizard before, this step is skipped. If you have not run the Prepare IBM MQ wizard before, the Prepare IBM MQ Wizard window displays a progress bar with the following message:

Status: Setting up IBM MQ Configuration

Wait until the progress bar completes.

3. The Prepare IBM MQ Wizard window displays a progress bar with the following message:

Status: Setting up the IBM MQ Service.

Wait until the progress bar completes.

4. IBM MQ attempts to detect whether you must configure IBM MQ for use with Windows Active Directory Server or later domain users. Depending on the results of the detection, IBM MQ does one of the following things:

- If IBM MQ detects that you need to configure IBM MQ for Windows Active Directory Server or later domain users, the Prepare IBM MQ Wizard window displays a message that starts:

IBM MQ does not have the authority to query information about your user account

Optionally, to see online help about configuring the domain account, select **More Information**. When you are finished, close the IBM MQ Help Center window to return to the current window.

Click **Next**, and go to step 5.

- If you are not installing on a Windows Active Directory Server or later domain server and IBM MQ cannot detect whether you need to configure IBM MQ for Windows Active Directory Server or later domain users, the Prepare IBM MQ Wizard window displays the following message:

Are any of the domain controllers in your network running Windows 2000 or later domain server?

If you select **Yes**, click **Next**, then go to step 5.

If you select **No**, click **Next**, then go to step 9.

If you select **Don't know**, you cannot continue. Select one of the other options, or click **Cancel** and contact your domain administrator.

- If IBM MQ detects that you do not need to configure IBM MQ for Windows Active Directory Server or later domain users, go to step 9.

Note: At any time, you can click **More Information** to view online help about configuring the domain account, or see “Configure IBM MQ accounts” on page 298. When you are finished, close the IBM MQ Help Center window to return to the current window.

5. The Prepare IBM MQ Wizard window displays the following message:

Do you need to configure IBM MQ for users defined on Windows 2000 or later domain controllers?

If you select **Yes**, click **Next**, then go to step 6.

If you select **No**, click **Next**, then go to step 9.

If you select **Don't know**, you cannot continue. Select one of the other options, or click **Cancel** and contact your domain administrator.

Note: At any time, you can click **More Information** to view online help about configuring the domain account, or see “Configure IBM MQ accounts” on page 298. When you are finished, close the IBM MQ Help Center window to return to the current window.

6. Give the domain user that you obtained from your domain administrator the access to run as a service.

- a. Click **Start > Run...**, type the command **secpol.msc** and click **OK**.
- b. Open **Security Settings > Local Policies > User Rights Assignments**. In the list of policies, right-click **Log on as a service > Properties**.
- c. Click **Add User or Group...** and type the name of the user you obtained from your domain administrator, and click **Check Names**.
- d. If prompted by a Windows Security window, type the user name and password of an account user or administrator with sufficient authority, and click **OK > Apply > OK**. Close the Local Security Policy window.

7. In the next window, enter the Domain and user ID of the domain user account that you obtained from your domain administrator. Either enter the Password for this account, or select the option **This account does not have a password**. Click **Next**.

8. The Prepare IBM MQ Wizard window displays a progress bar with the following message:

Status: Configuring IBM MQ with the special domain user account

Wait until the progress bar completes.

If there are any problems with the domain user account, a further window is displayed. Follow the advice on this window before you continue with this procedure.

9. The Prepare IBM MQ Wizard window displays a progress bar with the following message:

Status: Starting IBM MQ services

Wait until the progress bar completes.

10. Next, select the options that you require. The Prepare IBM MQ Wizard window displays the following message:

You have completed the Prepare IBM MQ Wizard

Select the options that you require, then click **Finish**. Select one or more from:

- **Remove the shortcut to this wizard from the desktop**

This option is available only if you have previously attempted installation, but you canceled the procedure from the Prepare IBM MQ wizard and you created a desktop shortcut to this wizard. Select this option to remove the shortcut. You do not need it now that you have completed the Prepare IBM MQ wizard.

- **Launch IBM MQ Explorer**

The IBM MQ Explorer allows you to view and administer your IBM MQ network.

- **Launch Notepad to view the release notes**

The release notes contain information about installing IBM MQ and also late-breaking news that is available after the published documentation is produced.

11. Follow the procedure described in "Checking for problems after installing."

Related information:

User rights required for an IBM MQ Windows Service

Checking for problems after installing:

These are optional tasks that you can use to check the installation if you believe there was a problem, or to verify installation messages after an unattended (silent) installation for example.

About this task

Use these steps as a guide to check the following files for messages:

Procedure

1. MSI *nnnnn*.LOG. This file is in your user Temp folder. It is an application log that contains English messages written during installation. The log includes a message indicating whether the installation was successful and complete.

This file is created if you have set up default logging.

2. If you used the launchpad to install IBM MQ, check MQv7_Install_YYYY-MM-DDTHH-MM-SS.log in your user Temp folder, where:

YYYY This is the year that you installed IBM WebSphere MQ Version 7.0

MM This is the month that you installed IBM MQ, for example this would be 09 if you installed in September

DD This is the day that you installed IBM MQ

HH-MM-SS

This is the time at which IBM MQ was installed

You can get to your user Temp directory by entering the following command at the command prompt:

```
cd %TEMP%
```

3. amqmjpse.txt. This file is in the IBM MQ data files folder (default C:\ProgramData\IBM\MQ). It is an application log that contains English messages written during installation by the Prepare IBM MQ wizard.

What to do next

1. Verify your installation, as described in “Verifying a server installation” on page 363

Configure IBM MQ accounts:

The IBM MQ service and queue managers check that any users attempting to access queue managers or queue manager resources such as queues, have the permission to access them.

Most networked Windows systems are members of a Windows domain where user accounts, other security principals, and security groups are maintained and managed by a directory service, Active Directory, running on a number of domain controllers. IBM MQ checks that only authorized users can access queue managers or queues.

In such networks, IBM MQ queue manager processes access the Active Directory information to find the security group membership of any users attempting to use IBM MQ resources. The accounts under which IBM MQ services run must be authorized to look up such information from the directory. In most Windows domains, local accounts defined at individual Windows servers cannot access directory information, so the IBM MQ services must run under a domain account that has the appropriate permission.

If the Windows server is not a member of a Windows domain or the domain has a reduced security or functional level, then the IBM MQ services can run under a local account that was created during installation.

Assuming that a domain account is needed, provide the information described in the Information for domain administrator to your domain administrator, and ask for one of the special accounts it describes. When you install the product, towards the end of the installation procedure, in the Prepare IBM MQ wizard, you are asked to enter details of this account (domain, user name, and password).

If a domain account is needed and you install IBM MQ without a special account (or without entering its details), many or all parts of IBM MQ do not work, depending upon the particular user accounts involved. Also, IBM MQ connections to queue managers that run under domain accounts on other systems might fail. The account can be changed by running the Prepare IBM MQ wizard and specifying the details of the account to be used.

For information about the user rights required to take advantage of the Active Directory support, see Using Active directory (Windows only).

For information about the user rights required to take advantage of the Kerberos authentication support, see Security.

Information for domain administrators:

Use this topic to understand how IBM MQ services check the authorization of user accounts attempting to access IBM MQ.

The user account must either have an individual IBM MQ authorization set or belong to a local group that has been authorized. A domain account can also be authorized through membership of a domain group included under an authorized local group through a single level of nesting.

The account under which the IBM MQ services are run must have the ability to query group memberships of domain accounts and have the authority to administer IBM MQ. Without the ability to query group memberships the access checks made by the services fail.

On most Windows domains, with domain controllers running Windows Active Directory, local accounts do not have the required authorization and a special domain user account with the required permissions must be used. The IBM MQ installer must be given the userid and password details so that they can be used to configure the IBM MQ service after the product is installed.

Typically, this special account has the IBM MQ administrator rights through membership of the domain group DOMAIN\Domain mqm. The domain group is automatically nested by the installation program under the local mqm group of the system on which IBM MQ is being installed.

See “Creating and setting up domain accounts for IBM MQ” for instructions on creating a suitable domain account.

Note: If an installer configures IBM MQ without a special account, many or all parts of IBM MQ do not work, depending upon the particular user accounts involved, as follows:

- An installer currently logged on with a domain user account is not be able to complete the Default Configuration, and the Postcard application does not work.
- IBM MQ connections to queue managers running under domain accounts on other systems might fail.
- Typical errors include “AMQ8066: Local mqm group not found” and “AMQ8079: Access was denied when attempting to retrieve group membership information for user 'abc@xyz'”.

Creating and setting up domain accounts for IBM MQ:

The following information is intended at Domain Administrators. Use this information to create and setup domain accounts for IBM MQ.

About this task

Repeat Steps 1 and 2 on page 301 for each domain that has user names that will install IBM MQ, to create an account for IBM MQ on each domain:

Procedure

1. Create a domain group with a special name that is known to IBM MQ and give members of this group the authority to query the group membership of any account:
 - a. Log on to the domain controller as an account with domain administrator authority.
 - b. From the Start menu, open Active Directory Users and Computers.
 - c. Find the domain name in the navigation pane, right-click it and select **New Group**.
 - d. Type a group name into the **Group name** field.

Note: The preferred group name is Domain mqm. Type it exactly as shown.

- Calling the group `Domain mqm` modifies the behavior of the “Prepare IBM MQ” wizard on a domain workstation or server. It causes the “Prepare IBM MQ” wizard automatically to add the group `Domain mqm` to the local `mqm` group on each new installation of IBM MQ in the domain.
 - You can install workstations or servers in a domain with no `Domain mqm` global group. If you do so, you must define a group with the same properties as `Domain mqm` group. You must make that group, or the users that are members of it, members of the local `mqm` group wherever IBM MQ is installed in a domain. You can place domain users into multiple groups. Create multiple domain groups, each group corresponding to a set of installations that you want to manage separately. Split domain users, according to the installations they manage, into different domain groups. Add each domain group or groups to the local `mqm` group of different IBM MQ installations. Only domain users in the domain groups that are members of a specific local `mqm` group can create, administer, and run queue managers for that installation.
 - The domain user that you nominate when installing IBM MQ on a workstation or server in a domain must be a member of the `Domain mqm` group, or of an alternative group you defined with same properties as the `Domain mqm` group.
- e. Leave **Global** clicked as the **Group scope**, or change it to **Universal**. Leave **Security** clicked as the **Group type**. Click **OK**.
- f. Follow these steps to assign permissions to the group based on the Windows version of the domain controller:

On Windows Server 2012 and Windows Server 2012 R2

- 1) In the Server Manager, click **Tools** then select **Active Directory Users and Computers** from the list box.
- 2) Select **View > Advanced Features**
- 3) Expand your domain name, then click **Users**
- 4) In the Users window, right-click **Domain mqm > Properties**
- 5) Click **Security > Advanced > Add...**
- 6) Click **Select principle**, then type `Domain mqm` and click **Check names > OK**
The **Name** field is prefilled with the string `Domain mqm (domain name\Domain mqm)`
- 7) In the **Applies to** list, select **Descendant User Objects**
- 8) In the **Permissions** list, select the **Read group membership** and **Read groupMembershipSAM** check boxes.
- 9) Click **OK > Apply > OK > OK**.

On Windows Server 2008 and Windows 2008 R2:

- 1) In the Server Manager navigation tree, click **Users**.
- 2) In the Server Manager action bar, click **View > Advanced features**
- 3) In the Users window, right-click **Domain mqm > Properties**
- 4) Click **Security > Advanced > Add**, then type `Domain mqm` and click **Check names > OK**
The **Name** field is prefilled with the string `Domain mqm (domain name\Domain mqm)`
- 5) Click **Properties**. In the **Apply to** list, select **Descendant User Objects**
- 6) In the **Permissions** list, select the **Read group membership** and **Read groupMembershipSAM** check boxes.
- 7) Click **OK > Apply > OK > OK**.

On Windows Server 2003:

- 1) In the Server Manager action bar, click **View > Advanced features > Active Directory Users and Computers**.
- 2) In the Server Manager navigation tree, search for the domain name. Select the domain name, right-click and select **Properties**.
- 3) Click **Security > Advanced > Add**. Type `Domain mqm` and click **Check names > OK**.
- 4) Click **Properties**. In the **Apply to** list, select **User Objects**

- 5) In the **Permissions** list, select the **Read group membership** and **Read groupMembershipSAM** check boxes.
- 6) Click **OK > Apply > OK > OK**.

On Windows Server 2000:

- 1) In the Server Manager navigation tree, search for the domain name. Select the domain name, right-click and select **Delegate Control Next**.
- 2) Click **Selected Groups and Users > Add...** Select Domain mqm and click **Add > OK**.
- 3) Select Domain mqm and click **Next**.
- 4) Click **Create a custom task to delegate** and click **Next**.
- 5) Select **Only the following objects in the folder**, and then check User Objects in the alphabetical list. Click **Next**.
- 6) Check **Property-specific**, then select the following options from the list:
 - **Read group membership**
 - **Read groupMembershipSAM**

Note: The list is in alphabetical order by the second word.

- 7) Click **OK** to close each window.
2. Create one or more accounts, and add them to the group:
 - a. In **Active Directory Users and Computers**, create a user account with a name of your choosing and add it to group Domain mqm (or a group that is a member of the local mqm group).
 - b. Repeat for all the accounts you want to create.
 3. Repeat Steps 1 on page 299 and 2 for each domain that has user names that will install IBM MQ, to create an account for IBM MQ on each domain.
 4. Use the accounts to configure each installation of IBM MQ:
 - a. Either use the same domain user account (as created in Step 1 on page 299) for each installation of IBM MQ, or create a separate account for each one, adding each to the Domain mqm group (or a group that is a member of the local mqm group).
 - b. When you have created the account or accounts, give one to each person configuring an installation of IBM MQ. They must enter the account details (domain name, user name, and password) into the Prepare IBM MQ wizard. Give them the account that exists on the same domain as their installing userid.
 - c. When you install IBM MQ on any system on the domain, the IBM MQ install program detects the existence of the Domain mqm group on the LAN, and automatically adds it to the local mqm group. (The local mqm group is created during installation; all user accounts in it have authority to manage IBM MQ). Thus all members of the "Domain mqm" group will have authority to manage IBM MQ on this system.
 - d. However, you do still need to provide a domain user account (as created in Step 1 on page 299) for each installation, and configure IBM MQ to use it when making its queries. The account details must be entered into the Prepare IBM MQ wizard that runs automatically at the end of installation (the wizard can also be run at any time from the **start** menu).
 5. Set the password expiry periods:
 - If you use just one account for all users of IBM MQ, consider making the password of the account never expire, otherwise all instances of IBM MQ will stop working at the same time when the password expires.
 - If you give each user of IBM MQ their own user account you will have more user accounts to create and manage, but only one instance of IBM MQ will stop working at a time when the password expires.

If you set the password to expire, warn the users that they will see a message from IBM MQ each time it expires - the message warns that the password has expired, and describes how to reset it.

6. Running IBM MQ as a service. If you need to run IBM MQ as a service, and then give the domain user (that you obtained from your domain administrator) the access to run as a service, carry out the following procedure:
 - a. Click **Start > Run...** Type the command `secpol.msc` and click **OK**.
 - b. Open **Security Settings > Local Policies > User Rights Assignments**. In the list of policies, right-click **Log on as a service > Properties**.
 - c. Click **Add User or Group...** Type the name of the user you obtained from your domain administrator, and click **Check Names**
 - d. If prompted by a Windows Security window, type the user name and password of an account user or administrator with sufficient authority, and click **OK > Apply > OK**. Close the Local Security Policy window.

Note: On Windows Server 2008 and Windows Server 2012 the User Account Control (UAC) is enabled by default.

The UAC feature restricts the actions users can perform on certain operating system facilities, even if they are members of the Administrators group. You must take appropriate steps to overcome this restriction.

Using the Default Configuration wizard:

About this task

You can use the Default Configuration wizard to add the first configured queue manager to this system. This enables you to connect easily with other queue managers in the same IBM MQ cluster. You can use the Default Configuration wizard to create, view, or alter your default configuration. You can also use this wizard to alter or display details of an existing queue manager that was created by the default configuration.

For a new installation of IBM MQ, creating a default configuration enables you to explore features of IBM MQ using the Postcard application, and the IBM MQ Explorer.

The Postcard application provides a fast and simple way to verify that your IBM MQ installation completed successfully. It uses the default queue manager that is created during the default configuration. If you want to use the Postcard application for verification, and you do not have any existing queue managers, run the Default Configuration wizard first.

If you have migrated existing queue managers, or created any queue managers since installing IBM MQ, you might not want to run the Default Configuration wizard. This is because you cannot create the default configuration if other queue managers already exist. If you have previously created any other queue managers on this system and you still want to set up a default configuration, you must delete them before you run the Default Configuration wizard.

Start the Default Configuration wizard by selecting **Create the Default Configuration** on the Welcome to IBM MQ Explorer Content view page.

Using the Welcome to IBM MQ Explorer Content view page:

The Welcome to IBM MQ Explorer Content view page points you to any relevant applications, documentation, tutorials, and education. This page is displayed the first time you launch IBM MQ Explorer.

You can use the items in the Welcome to IBM MQ Explorer Content view page to explore the facilities in IBM MQ. This page is launched the first time the IBM MQ Explorer is launched. The Welcome page can be viewed at any time from the Explorer by clicking **IBM MQ** in the Navigator view. There are links to the following subjects from this page:

Create the Default Configuration

Allows you to add a configured queue manager to this system for connecting easily with other queue managers in the same IBM MQ cluster. You can also use it to alter or display details of an existing queue manager created by the default configuration. This feature is available only using TCP/IP.

Note: If you migrated existing queue managers, or if you have created any queue managers after you installed IBM MQ, you might not want to use this facility. This is because you can only set up a default configuration if there are no queue managers already, and you might not want to delete your existing queue managers.

Launch Postcard

Allows you to try out IBM MQ messaging quickly and easily. You can send a message either to your own machine or to another named user's machine. It is described in detail in "Verify the installation using the Postcard application" on page 370.

Using the Help Center:

The Help Center gives you access to all task-oriented help, information on the IBM website, and a link to the IBM MQ product documentation if you have installed it from the IBM MQ Documentation CD.

The IBM MQ Help Center can be accessed from the IBM MQ Explorer by selecting **Help > Help Contents**.

Installing IBM MQ server on IBM i



Install IBM MQ for IBM i by installing the IBM MQ server in its primary language, installing samples and installing additional languages.

Before you begin

Note: Installing the latest version of the IBM MQ server includes client capabilities. Only install the stand-alone client if you do not need the server capabilities.

You have completed planning the installation, obtained the installation CDs and set the system values, see "Operating System configuration and tuning for IBM MQ on IBM i" on page 247. For a complete list of IBM MQ installable services and components for IBM i systems, see Installable services and components for IBM i

About this task

How to install the base IBM MQ server in its primary language, install samples and install translated versions from a choice of national-languages.

You can install only one instance of IBM MQ for IBM i in each partition of your server.

Procedure

1. Sign on to the system with a user profile that has *ALLOBJ special authority, for example QSECOFR.
2. Install IBM MQ for IBM i, V8.0 base product, and primary language.

```
RSTLICPGM LICPGM ( 5724H72 ) DEV ( install device ) OPTION ( *BASE ) OUTPUT ( *PRINT )
```

where the parameters of RSTLICPGM are,

LICPGM (5724H72)

The product identifier for IBM MQ for IBM i.

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION (*BASE)

Install the base IBM MQ for IBM i product.

Unspecified parameters

Unspecified parameters such as **RSTOBJ (*ALL)**, revert to defaults. The command installs both IBM MQ and the language files for the primary language of your system. For installing additional languages, see step 4.

3. Optional: Install the samples using the command:

```
RSTLICPGM LICPGM ( 5724H72 ) DEV ( install device ) OPTION ( 1 ) OUTPUT ( *PRINT )
```

Where the parameters of RSTLICPGM are,

LICPGM (5724H72)

The product identifier for IBM MQ for IBM i.

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION (1)

Install the samples for IBM MQ for IBM i.

OUTPUT (*PRINT)

The output is printed with the spooled output of the job.

4. Optional: To install additional languages, sign on to the system with a user profile that has *ALLOBJ special authority. Choose a language code from the table.

Table 53. Globalizations of IBM MQ for IBM i.

Language ID	Language
2909	Belgian English
2966	Belgian French MNCS (Multi-National Character Set)
2981	Canadian French MNCS
2975	Czech
2950	English uppercase
2924	English uppercase and lowercase
2984	English US DBCS
2938	English US uppercase DBCS

Table 53. Globalizations of IBM MQ for IBM i. (continued)

Language ID	Language
2928	French
2940	French MNCS
2929	German
2939	German MNCS
2976	Hungarian
2932	Italian
2942	Italian MNCS
2962	Japanese
2986	Korean
2978	Polish
2979	Russian
2989	Simplified Chinese
2931	Spanish

- If installing Japanese language feature code 2962, ensure the CCSID of the job installing the product is set to 939 and not 930. Do this to avoid problems with invariant lowercase characters in CCSID 930
CHGJOB CCSID(939)
- If the language feature code is not in the table then the product has not been translated into your language. You must choose one of the available language feature codes and install that version instead. You must manually change the system library list to use IBM MQ in that language load.
CHGSYSLIBL LIB(QSYS2924)
- If you are using Korean DBCS and you configure your terminal emulators to 24*80 sessions you might find that EDTF incorrectly displays DBCS characters in MQ error log messages that extend beyond 80 columns. To avoid this, configure your terminal emulators to use sessions capable of displaying 132 columns, for example 27*132.
- Issue the following command specifying the appropriate language ID:
RSTLICPGM LICPGM(5724H72) DEV(*install device*) RSTOBJ(*LNG) LNG(*language ID*)

This installs the commands, message file, and panel groups into the relevant QSYS library for the language. For example, library QSYS2928 is used for French. If this QSYS29nn library does not exist, it is created by the RSTLICPGM command.

5. To ensure that the product has loaded correctly, issue the Display Software Resources (DSPSFWRSC) command and check that the licensed program 5724H72 is listed. If you have installed the base and the optional samples, you see:

```
Resource
ID   Option Feature Description
5724H72 *BASE 5050 IBM MQ for IBM i
5724H72 *BASE 2924 IBM MQ for IBM i
5724H72 1 5050 IBM MQ for IBM i - Samples
```

6. Press F11, while viewing the Display Software Resources screen, and you see the library and version number of the products installed:

Resource		Feature			
ID	Option	Feature	Type	Library	Release
5724H72	*BASE	5050	*CODE	QMQM	V8R0M0
5724H72	*BASE	2924	*LNG	QMQM	V8R0M0
5724H72	1	5050	*CODE	QMOMSAMP	V8R0M0

7. If you have installed additional language versions, you also see entries for these versions. For example, if you have installed the French version, for which the language ID is 2928, you see:
- a.

Resource		Feature		
ID	Option	Feature	Description	
5724H72	*BASE	2928	IBM MQ for IBM i	

- b. and when you press F11:

Resource		Feature			
ID	Option	Feature	Type	Library	Release
5724H72	*BASE	2928	*LNG	QSYS2928	V8R0M0

8. Use the command DSPMQMVER to check exactly what version you have installed. For V8R0M0, it reports:

Version: 8.0.0.0

9. Do the post installation tasks of checking for updates, checking program authorities and starting the IBM MQ subsystem, see “Post installation tasks for IBM MQ on IBM i” on page 314.

What to do next

If you want to see how the installation went in more detail, perform one or more of the following tasks:

- View the log file using the DSPJOBLOG command.
- View the spoolfile generated from the RSTLICPGM command.

If the installation of IBM MQ fails, see “Installation failures for IBM i” on page 316.

Related concepts:

“Uninstalling IBM MQ for IBM i” on page 400

There are two ways of uninstalling IBM MQ for IBM i.

Non-interactive installation of IBM MQ server on IBM i



You can perform a non-interactive installation of IBM MQ using the CALL PGM(QSYS/QLPACAGR) command. A non-interactive installation is also known as a silent, or unattended installation.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

This topic describes the non-interactive installation of a server.

Procedure

1. Pre-agree the license terms and conditions for the base by running the command,
`CALL PGM (QSYS/QLPACAGR) PARM ('5724H72' 'V8R0M0' '0000' 0)`

Where the parameters of **PARM** are,

5724H72

The product identifier for IBM MQ for IBM i.

V8R0M0

The version, release, and modification level.

0000

The option number for the IBM MQ product.

0 Unused error structure.

2. Optionally pre-agree the license terms and conditions for the samples by running the command,
`CALL PGM (QSYS/QLPACAGR) PARM ('5724H72' 'V8R0M0' '0001' 0)`

Where the parameters of **PARM** are,

5724H72

The product identifier for IBM MQ for IBM i.

V8R0M0

The version, release, and modification level.

0001

The option number for the IBM MQ product.

0 Unused error structure.

3. Install IBM MQ for IBM i, V8.0 base product, and primary language.

```
RSTLICPGM LICPGM ( 5724H72 ) DEV ( install device ) OPTION ( *BASE ) OUTPUT ( *PRINT )
```

where the parameters of RSTLICPGM are,

LICPGM (5724H72)

The product identifier for IBM MQ for IBM i.

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION (*BASE)

Install the base IBM MQ for IBM i product.

Unspecified parameters

Unspecified parameters such as **RSTOBJ (*ALL)**, revert to defaults. The command installs both IBM MQ and the language files for the primary language of your system. For installing additional languages, see step 4.

4. Optional: Install the samples using the command:

```
RSTLICPGM LICPGM ( 5724H72 ) DEV ( install device ) OPTION ( 1 ) OUTPUT ( *PRINT )
```

Where the parameters of RSTLICPGM are,

LICPGM (5724H72)

The product identifier for IBM MQ for IBM i.

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION (1)

Install the samples for IBM MQ for IBM i.

OUTPUT (*PRINT)

The output is printed with the spooled output of the job.

- Optional: To install additional languages, sign on to the system with a user profile that has *ALLOBJ special authority. Choose a language code from the table.

Table 54. Globalizations of IBM MQ for IBM i.

Language ID	Language
2909	Belgian English
2966	Belgian French MNCS (Multi-National Character Set)
2981	Canadian French MNCS
2975	Czech
2950	English uppercase
2924	English uppercase and lowercase
2984	English US DBCS
2938	English US uppercase DBCS
2928	French
2940	French MNCS
2929	German
2939	German MNCS
2976	Hungarian
2932	Italian
2942	Italian MNCS
2962	Japanese
2986	Korean
2978	Polish
2979	Russian
2989	Simplified Chinese
2931	Spanish

- If installing Japanese language feature code 2962, ensure the CCSID of the job installing the product is set to 939 and not 930. Do this to avoid problems with invariant lowercase characters in CCSID 930
CHGJOB CCSID(939)
- If the language feature code is not in the table then the product has not been translated into your language. You must choose one of the available language feature codes and install that version instead. You must manually change the system library list to use IBM MQ in that language load.
CHGSYSLIBL LIB(QSYS2924)
- If you are using Korean DBCS and you configure your terminal emulators to 24*80 sessions you might find that EDTF incorrectly displays DBCS characters in MQ error log messages that extend beyond 80 columns. To avoid this, configure your terminal emulators to use sessions capable of displaying 132 columns, for example 27*132.
- Issue the following command specifying the appropriate language ID:
RSTLICPGM LICPGM(5724H72) DEV(*install device*) RSTOBJ(*LNG) LNG(*language ID*)

This installs the commands, message file, and panel groups into the relevant QSYS library for the language. For example, library QSYS2928 is used for French. If this QSYS29nn library does not exist, it is created by the RSTLICPGM command.

- To ensure that the product has loaded correctly, issue the Display Software Resources (DSPSFWRSC) command and check that the licensed program 5724H72 is listed. If you have installed the base and the optional samples, you see:

```
Resource
ID   Option Feature Description
5724H72 *BASE 5050 IBM MQ for IBM i
5724H72 *BASE 2924 IBM MQ for IBM i
5724H72 1    5050 IBM MQ for IBM i - Samples
```

- Press F11, while viewing the Display Software Resources screen, and you see the library and version number of the products installed:

```
Resource      Feature
ID   Option Feature Type Library Release
5724H72 *BASE 5050 *CODE QMQM V8R0M0
5724H72 *BASE 2924 *LNG QMQM V8R0M0
5724H72 1    5050 *CODE QMQMSAMP V8R0M0
```

- If you have installed additional language versions, you also see entries for these versions. For example, if you have installed the French version, for which the language ID is 2928, you see:

a.

```
Resource
ID   Option Feature Description
5724H72 *BASE 2928 IBM MQ for IBM i
```

b. and when you press F11:

```
Resource      Feature
ID   Option Feature Type Library Release
5724H72 *BASE 2928 *LNG QSYS2928 V8R0M0
```

- Use the command DSPMQMVER to check exactly what version you have installed. For V8R0M0, it reports:

```
Version: 8.0.0.0
```

- Do the post installation tasks of checking for updates, checking program authorities and starting the IBM MQ subsystem, see “Post installation tasks for IBM MQ on IBM i” on page 314.

What to do next

If you want to see how the installation went in more detail, perform one or more of the following tasks:

- View the log file using the DSPJOBLOG command.
- View the spoolfile generated from the RSTLICPGM command.

If the installation of IBM MQ fails, see “Installation failures for IBM i” on page 316.

Installing IBM MQ Managed File Transfer on IBM i



Install IBM MQ Managed File Transfer for IBM i by installing IBM MQ Java Messaging and Web Services server in its primary language, and installing additional options.

Before you begin

Note: Installing the latest version of IBM MQ Managed File Transfer includes client capabilities.

You have completed planning the installation, obtained the installation CDs and set the system values, see “Operating System configuration and tuning for IBM MQ on IBM i” on page 247.

You have installed:

Table 55.

Program	Option	Description
5761JV1	14 or 15	Java SE 7 32 bit or Java SE 7 64 bit
5770SS1	39	International Components for Unicode
5724L26	*BASE	IBM MQ Java Messaging and Web Services

About this task

How to install base IBM MQ Managed File Transfer in its primary language, and install the other options.

You can install only one instance of IBM MQ Managed File Transfer for IBM i in each partition of your server.

Procedure

1. Sign on to the system with a user profile that has *ALLOBJ special authority, for example QSECOFR.
2. Install IBM MQ Managed File Transfer for IBM i, V8.0 base product.

```
RSTLICPGM LICPGM ( 5725M50 ) DEV ( install device ) OPTION ( *BASE ) OUTPUT ( *PRINT )
```

where the parameters of RSTLICPGM are,

LICPGM (5725M50)

The product identifier for IBM MQ Managed File Transfer for IBM i.

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, 0PT01.

OPTION (*BASE)

Install the base IBM MQ Managed File Transfer for IBM i product.

Unspecified parameters

Unspecified parameters such as **RSTOBJ (*ALL)**, revert to defaults. The command installs both IBM MQ and the language files for the primary language of your system.

3. Optional: Install the tools using the command:

```
RSTLICPGM LICPGM ( 5725M50 ) DEV ( install device ) OPTION ( 2 ) OUTPUT ( *PRINT )
```

Where the parameters of RSTLICPGM are,

LICPGM (5725M50)

The product identifier for IBM MQ Managed File Transfer for IBM i.

DEV (install device)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION (2)

Install the tools for IBM MQ Managed File Transfer for IBM i.

OUTPUT (*PRINT)

The output is printed with the spooled output of the job.

Repeat step 3 on page 310 for options 3 (agent) and 4 (services)

- To ensure that the product has loaded correctly, issue the Display Software Resources (DSPSFWRSC) command and check that the licensed program 5725M50 is listed. If you have installed the base and the optional tools, you see:

Resource ID	Option	Feature	Description
5725M50	*BASE	5050	Managed File Transfer for IBMi
5725M50	*BASE	2924	Managed File Transfer for IBMi
5725M50	2	5050	Managed File Transfer for IBMi - Tools

- Press F11, while viewing the Display Software Resources screen, and you see the library and version number of the installed products:

Resource ID	Option	Feature	Type	Library	Release
5725M50	*BASE	5050	*CODE	QMOMMFT	V8R0M0
5725M50	*BASE	2924	*LNG	QMOMMFT	V8R0M0
5725M50	2	5050	*CODE	MFTTOOL	V8R0M0

- Do the post installation tasks of checking for updates, checking program authorities, and starting the Managed File Transfer subsystem.

What to do next

If you want to see how the installation went in more detail, perform one or more of the following tasks:

- View the log file using the DSPJOBLOG command.
- View the spoolfile generated from the RSTLICPGM command.

If the installation of IBM MQ fails, see “Installation failures for IBM i” on page 316.

IBM MQ for IBM i Electronic Software Download Installation

You can perform an installation of IBM MQ for IBM i Version 8.0 from an installation image downloaded from IBM.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

Two installation images are provided as zip files, a client and server image. These images contain all the licensed programs, and a client only image for the clients only.

The client and server image contains all seven compressed IBM i save files (**SAVF**), while the client image contains four save files. The save files are:

- MQ80BASE - IBM MQ client & server base program objects
- MQ80SAMP - IBM MQ client & server samples
- MQ80EN24 - IBM MQ client & server English US (2924) language objects

plus the client only images:

- MQ80CBASE - IBM MQ client
- MQ80CSAMP - IBM MQ client samples
- MQ80JBASE - IBM MQ Java
- MQ80JSAMP - IBM MQ Java samples

Procedure

1. Download one of the installation images and unzip it to a temporary directory.
2. On IBM i, create a library containing sufficient empty save files to hold the uploaded files by using the commands:

```
CRTLIB LIB(MQ80PROD)
CRTSAVF FILE(MQ80PROD/MQ80BASE) /* Server and Client */
CRTSAVF FILE(MQ80PROD/MQ80SAMP) /* Server and Client Samples */
CRTSAVF FILE(MQ80PROD/MQ80EN24) /* 2924 English */
CRTSAVF FILE(MQ80PROD/MQ80CBASE) /* Standalone Client */
CRTSAVF FILE(MQ80PROD/MQ80CSAMP) /* Standalone Client Samples */
CRTSAVF FILE(MQ80PROD/MQ80JBASE) /* Java and JMS Classes */
CRTSAVF FILE(MQ80PROD/MQ80JSAMP) /* Java and JMS Samples */
```

For additional languages

```
CRTSAVF FILE(MQ80PROD/MQ80EN09) /* 2909 Belgian English */
CRTSAVF FILE(MQ80PROD/MQ80FR28) /* 2928 French */
CRTSAVF FILE(MQ80PROD/MQ80JA30) /* 2930 Japanese */
CRTSAVF FILE(MQ80PROD/MQ80ES31) /* 2931 Spanish */
CRTSAVF FILE(MQ80PROD/MQ80IT32) /* 2932 Italian */
CRTSAVF FILE(MQ80PROD/MQ80EN38) /* 2938 English DBCS UPPERCASE */
CRTSAVF FILE(MQ80PROD/MQ80FR40) /* 2940 French MNCS */
CRTSAVF FILE(MQ80PROD/MQ80IT42) /* 2942 Italian MNCS */
CRTSAVF FILE(MQ80PROD/MQ80FR66) /* 2966 French MNCS */
CRTSAVF FILE(MQ80PROD/MQ80FR81) /* 2981 French MNCS */
CRTSAVF FILE(MQ80PROD/MQ80EN84) /* 2984 English DBCS */
CRTSAVF FILE(MQ80PROD/MQ80CZ75) /* 2975 Czech */
CRTSAVF FILE(MQ80PROD/MQ80HU76) /* 2976 Hungarian */
CRTSAVF FILE(MQ80PROD/MQ80PL78) /* 2978 Polish */
CRTSAVF FILE(MQ80PROD/MQ80RU79) /* 2979 Russian */
CRTSAVF FILE(MQ80PROD/MQ80PT80) /* 2980 Portugese/Brazilian */
CRTSAVF FILE(MQ80PROD/MQ80JA62) /* 2962 Japanese */
CRTSAVF FILE(MQ80PROD/MQ80K086) /* 2986 Korean */
CRTSAVF FILE(MQ80PROD/MQ80ZH89) /* 2989 Chinese */
CRTSAVF FILE(MQ80PROD/MQ80DE29) /* 2929 German */
CRTSAVF FILE(MQ80PROD/MQ80DE39) /* 2939 German */
```

3. Start an ftp session to your IBM i machine and upload the required save files with the commands:

```
ftp (your_ibmi_hostname)
bin
put MQ80BASE MQ80PROD/MQ80BASE
put MQ80SAMP MQ80PROD/MQ80SAMP
put MQ80EN24 MQ80PROD/MQ80EN24
put MQ80CBASE MQ80PROD/MQ80CBASE
put MQ80CSAMP MQ80PROD/MQ80CSAMP
put MQ80JBASE MQ80PROD/MQ80JBASE
put MQ80JSAMP MQ80PROD/MQ80JSAMP
```

For additional language loads:


```

put MQ80EN09 MQ80PROD/MQ80EN09
put MQ80FR28 MQ80PROD/MQ80FR28
put MQ80JA30 MQ80PROD/MQ80JA30
put MQ80ES31 MQ80PROD/MQ80ES31
put MQ80IT32 MQ80PROD/MQ80IT32
put MQ80EN38 MQ80PROD/MQ80EN38
put MQ80FR40 MQ80PROD/MQ80FR40
put MQ80IT42 MQ80PROD/MQ80IT42
put MQ80FR66 MQ80PROD/MQ80FR66
put MQ80FR81 MQ80PROD/MQ80FR81
put MQ80EN84 MQ80PROD/MQ80EN84
put MQ80CZ75 MQ80PROD/MQ80CZ75
put MQ80HU76 MQ80PROD/MQ80HU76
put MQ80PL78 MQ80PROD/MQ80PL78
put MQ80RU79 MQ80PROD/MQ80RU79
put MQ80PT80 MQ80PROD/MQ80PT80
put MQ80JA62 MQ80PROD/MQ80JA62
put MQ80K086 MQ80PROD/MQ80K086
put MQ80ZH89 MQ80PROD/MQ80ZH89
put MQ80DE29 MQ80PROD/MQ80DE29
put MQ80DE39 MQ80PROD/MQ80DE39

```

4. To prepare for installation of IBM MQ for IBM i, sign on to your IBM i machine and ensure that you have followed the instructions detailed in “Preparing the system” on page 233
5. Enter the **RSTLICPGM** commands, specifying the install device as *SAVF and naming the save file containing the options that you want to install. The IBM MQ Java licensed program can be installed standalone or can coexist with any of the other licensed programs.

The IBM MQ client can be installed standalone, but it can only coexist with the IBM MQ Java on the same system.

Attempting to install the IBM MQ server on a system where the IBM MQ client is already installed performs a slip install upgrade, replacing the client with the server licensed program.

Attempting to install the IBM MQ client standalone over the top of an existing server licensed program is not possible, and the install fails.

For example:

```

/* IBM MQ Client and Server program objects */
RSTLICPGM LICPGM(5724H72) DEV(*SAVF) SAVF(MQ80PROD/MQ80BASE) +
RSTOBJ(*PGM) OPTION(*BASE) OUTPUT(*PRINT)

/* IBM MQ Client & Server English 2924 Language Load */
RSTLICPGM LICPGM(5724H72) DEV(*SAVF) SAVF(MQ80PROD/MQ80EN24) +
RSTOBJ(*LNG) LNG(2924) OUTPUT(*PRINT)

/* Additional languages - alter SAVF and LNG parameters... */
/* IBM MQ Client & Server Japanese 2930 Language Load */
RSTLICPGM LICPGM(5724H72) DEV(*SAVF) SAVF(MQ80PROD/MQ80JA30) +
RSTOBJ(*LNG) LNG(2930) OUTPUT(*PRINT)

/* IBM MQ Client & Server Samples */
RSTLICPGM LICPGM(5724H72) DEV(*SAVF) SAVF(MQ80PROD/MQ80SAMP) +
OPTION(1) OUTPUT(*PRINT)

/* IBM MQ Java */
RSTLICPGM LICPGM(5724L26) DEV(*SAVF) SAVF(MQ80PROD/MQ80JBASE) +
OPTION(*BASE) OUTPUT(*PRINT)

/* IBM MQ Java Samples */
RSTLICPGM LICPGM(5724L26) DEV(*SAVF) SAVF(MQ80PROD/MQ80JSAMP) +
OPTION(1) OUTPUT(*PRINT)

/* IBM MQ Client */
RSTLICPGM LICPGM(5725A49) DEV(*SAVF) SAVF(MQ80PROD/MQ80CBASE) +
OPTION(*BASE) OUTPUT(*PRINT)

/* IBM MQ Client Samples */
RSTLICPGM LICPGM(5725A49) DEV(*SAVF) SAVF(MQ80PROD/MQ80CSAMP) +
OPTION(1) OUTPUT(*PRINT)

```

6. Do the post installation tasks of checking for updates, checking program authorities and starting the IBM MQ subsystem, see “Post installation tasks for IBM MQ on IBM i.”

What to do next

If you want to see how the installation went in more detail, perform one or more of the following tasks:

- View the log file using the DSPJOBLOG command.
- View the spoolfile generated from the RSTLICPGM command.

If the installation of IBM MQ fails, see “Installation failures for IBM i” on page 316.

Post installation tasks for IBM MQ on IBM i



Tasks to perform after you have installed IBM MQ for IBM i, and before using it.

About this task

When you have correctly installed IBM MQ for IBM i on your system:

Procedure

1. See the IBM MQ website at: <http://www.ibm.com/software/integration/wmq/index.html> for the latest product information.
2. Install and apply all fix packs.
3. Where you have more than one system and a mixture of releases of OS/400 or IBM i, and IBM MQ, you must take care when compiling CL programs. You must compile CL programs either on the system they are to run on, or on one with an identical combination of releases of OS/400 or IBM i,

and IBM MQ. When you install later versions of IBM MQ, delete all IBM MQ commands from previous releases in any QSYSVvRrMm libraries using the QSYS/DLTCMD command.

4. If you have not installed IBM MQ on your system before, you must add user profiles to the QMQMADM group profile. Make all user profiles that are to be used for creating and administering queue managers members of the QMQMADM group profile, using the command CHGUSRPRF.
 - a. Start the IBM MQ subsystem, by issuing the command:
STRSBS SBS(DQMQM/QMQM)

Note: The subsystem must be started after each IPL of the system, so you might choose to start it as part of your system startup process.

5. Create the system-default objects. The system-default objects are created automatically when you issue the CRTMQM command to create a queue manager. For example: CRTMQM MQMNAME(QMGRNAME) ASP(*SYSTEM). You can refresh them using the STRMQM command (Warning: this command will replace any existing default objects). For example: STRMQM MQMNAME(QMGRNAME) RDEFSYS(*YES). Refer to the onscreen help for information about using this command.

Note: on the command STRMQM MQMNAME(QMGRNAME) RDEFSYS(*YES):

- The command does not recreate the objects, it performs a CRTxxxx REPLACE(*YES) for all of the SYSTEM.* objects.
- This means that it refreshes the parameters on the objects back to their defaults. So if, for example, on the SYSTEM.DEFAULT.LOCAL.QUEUE object, TRGENBL had previously been changed to *YES, then, when the command is run, it is changed back to TRGENBL(*NO).
- If any messages exist on a queue, they are not removed, because the queues are not physically deleted.
- The contents of the SYSTEM.AUTH.DATA.QUEUE are untouched when this command is run.
- So, if the contents of this (or any other significant queue) become corrupt, it must be physically deleted and recreated either from scratch, or from a backup.

Results

You are now ready to start using IBM MQ for IBM i.

Note: When you install IBM MQ for IBM i, two user profiles are created:

- QMQM
- QMQMADM

These two objects are central to the correct running of IBM MQ for IBM i. Do not alter or delete them. If you do, IBM cannot guarantee correct behavior of your product.

If you uninstall IBM MQ and data, these profiles are deleted. If you uninstall IBM MQ only, these profiles are retained.

Installation failures for IBM i



If the installation of IBM MQ Server or Client for IBM i fails, remove the installed and partially installed objects before attempting reinstallation:

Procedure

1. Delete installed options using `DLTLICPGM LICPGM(5725A49)OPTION(*ALL)`.
2. Delete partially installed options by deleting the QMQM library (and the QMQMSAMP libraries if necessary).
3. Delete the IFS directory `/QIBM/ProdData/mqm` and its subdirectories using the `EDTF` command, for example: `EDTF STMF ('/QIBM/ProdData')` and select **option 9** for the `mqm` directory.

If the installation of IBM MQ Java fails, remove the partly installed objects before attempting reinstallation:

- a. Delete the QMQMJAVA library.
- b. Delete the IFS directory `/QIBM/ProdData/mqm/java` and its subdirectories using the `EDTF` command, for example:

```
EDTF STMF ( '/QIBM/ProdData/mqm' )
```

Select option 9 against the Java directory.

Converting a trial license on UNIX, Linux, and Windows

Convert a trial license to a full license without reinstalling IBM MQ.

When the trial license expires, the “count-down” displayed by the `strmqm` command informs you the license has expired, and the command does not run.

Before you begin

1. IBM MQ is installed with a trial license.
2. You have access to the installation media of a fully licensed copy of IBM MQ.

About this task

Run the `setmqprd` command to convert a trial license to a full license.

If you do not want to apply a full license to your trial copy of IBM MQ, you can uninstall it at any time.

Procedure

1. Obtain the full license from the fully licensed installation media.

The full license file is **amqpcert.lic**:

- On UNIX and Linux, it is in the `/MediaRoot/licenses` directory on the installation media.
- On Windows it is in the `\MediaRoot\licenses` directory on the installation media. It is installed into the `bin` directory on the IBM MQ installation path.
- On IBM i, issue the command

```
CALL PGM(QMQM/SETMQPRD) PARM('/QOPT/OPT01/amqpcert.lic')
```

2. Run the `setmqprd` command from the installation that you are upgrading:

```
$MQ_INSTALLATION_PATH/bin/setmqprd /MediaRoot/licenses/amqpcert.lic
```

Related information:

setmqprd

Displaying messages in your national language on UNIX and Linux systems

To display messages from a different national language message catalog, you must install the appropriate catalog and set the **LANG** environment variable.

About this task

Non-AIX platforms

Messages in U.S. English are automatically installed with IBM MQ.

AIX Messages in the language specified by the locale selected on your machine at install time are installed by default.

To find out which language is currently in use, run the **locale** command.

If this returns a language which is not one of the national languages provided by IBM MQ, you must select a national language, otherwise you will not get a message catalog installed on your system.

Message catalogs for all languages are installed in *MQ_INSTALLATION_PATH/msg/language identifier*, where *language identifier* is one of the identifiers in Table 56.

If you require messages in a different language, perform the following steps:

Procedure

1. Install the appropriate message catalog (see “Choosing what to install” on page 194).
2. To select messages in a different language, ensure the **LANG** environment variable is set to the identifier for the language you want to install:

Table 56. Language identifiers

Identifier	Language
cs_CZ	Czech
de_DE	German
es_ES	Spanish
fr_FR	French
hu_HU	Hungarian
it_IT	Italian
ja_JP	Japanese
ko_KR	Korean
pl_PL	Polish
pt_BR	Brazilian Portuguese
ru_RU	Russian
zh_CN	Simplified Chinese
zh_TW	Traditional Chinese

AIX has some additional message catalogs:

Table 57. AIX specific language identifiers

Identifier	Language
Ja_JP	Japanese
Zh_CN	Simplified Chinese
Zh_TW	Traditional Chinese

Displaying messages in your national language on Windows systems

Windows

To display messages from a different national language message catalog, you must either set the **MQS_FORCE_NTLANGID** environment variable, or change a Regional setting.

About this task

Messages in U.S. English are automatically installed with IBM MQ

Messages in the national languages that IBM MQ supports are automatically installed. Messages are displayed in the national language, based on the following order:

1. The value of the **MQS_FORCE_NTLANGID** environment variable, if set.
2. The Region Format of the user that is displaying the message, if the language specified by the Region Format is supported by IBM MQ.
3. The Administrative system locale if the language specified by the system locale is supported by IBM MQ.
4. US English, if no other supported language can be determined.

Note: The queue manager is usually launched by a service on the machine, and hence is running under its own user account (for example MUSR_MQADMIN) or a specific domain account provided during install time. See Security on Windows for more information.

If you require messages in a language other than the one associated with the Region Format of a user account, perform the following steps:

Procedure

1. Globally set the **MQS_FORCE_NTLANGID** environment variable, to the language identifier of the desired language, for messages displayed by the queue manager. You should set the **MQS_FORCE_NTLANGID** system wide. Otherwise, every user displaying messages needs to have the environment variable set individually.

The language identifier values, represented in hexadecimal notation, are listed in the following Microsoft document: [Language Identifier Constants and Strings](#)

2. Reboot machines where queue managers are running as a service, for the environment variable to take effect.

Installing an IBM MQ client

After preparing your system for installation, you can install an IBM MQ client by following the appropriate instructions for your platform. After installation, you might want to verify your installation to check that installation has been successful.

Before you start the installation procedure, make sure that you have prepared your system as described in *Preparing the system*

To begin the installing procedure, select the appropriate platform:

- “Installing an IBM MQ client on AIX systems”
- “Installing an IBM MQ client on HP Integrity NonStop Server systems” on page 321
- “Installing an IBM MQ client on HP-UX systems” on page 323
- “Installing an IBM MQ client on Linux” on page 325
- “Installing an IBM MQ client on Solaris” on page 333
- “Installing an IBM MQ client on Windows systems” on page 337
- “Installing an IBM MQ client on IBM i” on page 349

Installing IBM MQ clients and servers on the same system

To install an IBM MQ client on a system that is already running an IBM MQ server, use the appropriate Server DVD. Use a Client DVD to install an IBM MQ client only on a system that is not running an IBM MQ server.

If you install an IBM MQ client from a Client DVD and later decide to install the IBM MQ server on the same system, you must first remove all the client components from the system. Then use the appropriate Server DVD to install both the server and client components. You cannot install an IBM MQ server on a system that already has client components installed from a Client DVD.

Remember that even if your client and server are installed on the same system, you must still define the MQI channel between them. See *Defining MQI channels* for details.

Installing an IBM MQ client on AIX systems

You can interactively install the IBM MQ client for AIX using `smit`.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

IBM MQ is supplied as a set of filesets that are installed using the standard AIX installation tools. The procedure uses the System Management Interface Tool (`smit`), but you can choose to use `installp`, `geninstall` or the web-based System Manager. You can select which components you want to install. The components and filesets are listed in “Choosing what to install” on page 194. You must install at least the Runtime and Client components.

This procedure installs IBM MQ into the default location. If you want to install to a non-default location, you must use `installp`, see “Non-interactive installation of IBM MQ client for AIX” on page 320.

Procedure

1. Log in as root, or switch to the superuser using the `su` command.

2. Make your current directory the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Select the required **smit** window using the following sequence:
 - Software Installation and Maintenance
 - Install and Update Software
 - Install and Update from ALL Available Software
4. Click **List** to display the input device or directory for the software and select the location that contains the installation images.
5. Select the **SOFTWARE to install** field to obtain a list of available filesets, and select the filesets you want to install. Ensure that you include the appropriate message catalog if you require messages in a language different from the language specified by the locale specified on your system. Enter **ALL** to install all applicable filesets.
6. Change **Preview new LICENSE agreements?** to **yes** and press Enter to view the license agreements.
7. If you have a previous version of the product on your system, change the **Automatically install requisite software** to **no**.
8. Change **ACCEPT new license agreements?** to **yes** and press Enter to accept the license agreements.
9. Change **Preview new LICENSE agreements?** to **no** and press Enter to install IBM MQ.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see *Changing the primary installation*.

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see *setmqenv* and *crtmqenv*.
- For instructions on how to verify your installation, see “Verifying a client installation” on page 374.

Related tasks:

“Uninstalling IBM MQ on AIX” on page 389

On AIX, you can uninstall the IBM MQ server or client using the System Management Interface Tool (SMIT) or the **installp** command.

Non-interactive installation of IBM MQ client for AIX

Silently install IBM MQ client from the command line using the AIX **installp** command.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

Note: Installation to a non-default location is *not* supported on systems that have the AIX Trusted Computing Base (TCB) enabled.

About this task

You can use this method to install to a non-default location, and can select which components you want to install. The components and filesets are listed in “Choosing what to install” on page 194. You must install at least the Runtime and Client components.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Set your current directory to the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Install the product in one of the following ways:

- Install the whole product in the default location:

```
installp -acgXYd . all
```
- Install selected filesets in the default location:

```
installp -acgXYd . list of file sets
```
- Install the whole product in a non-default location using the -R flag:

```
installp -R USIL_Directory -acgXYd . all
```
- Install selected filesets in a non-default location using the -R flag:

```
installp -R USIL_Directory -acgXYd . list of file sets
```

where the directory specified with the -R flag is an AIX User Specified Install Location (USIL) directory which exists before the command is run; it must not contain any spaces or `usr/mqm`.

IBM MQ is installed underneath the directory specified. For example, if `/USIL1` is specified, the IBM MQ product files are located in `/USIL1/usr/mqm`. This location is known as the `MQ_INSTALLATION_PATH`.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see “[Verifying a client installation](#)” on page 374.

Installing an IBM MQ client on HP Integrity NonStop Server systems

Installing an IBM MQ client on a HP Integrity NonStop Server system.

Before you begin

Before you start the installation procedure, make sure that you complete the necessary steps that are outlined in “[Setting up the user and group on HP Integrity NonStop Server](#)” on page 236.

About this task

After preparing your system for installation, install the IBM MQ client for HP Integrity NonStop Server by following the instructions. After installation, you might want to verify your installation to check that it installed successfully. There are three steps to the installation:

1. Downloading the installation package.
2. Running the installer.
3. Setting the environment.

Procedure

1. Log in to the OSS user ID that owns the installation. The OSS user ID must have MQM as its primary group.

2. Download the installation package file. Ensure that you use "binary mode" when you download the installation package file to OSS. Any corruption in the file causes the self-extracting archive to fail to run. After you have downloaded the package file, ensure that it has read and execute permissions for the user ID that is installing the package.
3. Set the `_RLD_FIRST_LIB_PATH` variable to `<install path>/opt/mqm/bin`
4. Optional: Make your current directory the location of the installation file.
5. Type the following command to start the interactive installation procedure:
`./<name of package file> -i <OSS install_root> -g <Guardian install_root>`
 where
`<name of package file>` is the name of the installation package.
`<OSS install_root>` is the OSS root directory of the new installation.
`<Guardian install_root>` is the Guardian subvolume for the new installation.
 Both `-i` and `-g` options are mandatory.
 - `-i` specifies the new or empty OSS directory that contains the `opt/mqm` and `var/mqm` directories of the installation.
 - `-g` specifies the subvolume into which the Guardian components of the IBM MQ client on a HP Integrity NonStop Server are installed. The Guardian subvolume can be specified in either OSS-form or Guardian-form and can be abbreviated. The Guardian subvolume specification is not case sensitive. The following are examples of valid Guardian subvolume specifications:
 - `/G/vol/subvol`
 - `vol/subvol`
 - `\$VOL.SUBVOL`
 - `vol.subvol`
6. Optional: For OSS, set your environment by installing the binaries into your path. To do this, type the following command:
`export PATH=$PATH:<OSS_install_root>/opt/mqm/bin`
 where `<OSS_install_root>` is the OSS root directory of the new installation.

Example

To install the IBM MQ client for HP Integrity NonStop Server from package `mat1.run`, type the following command:

```
./mat1.run -i ~install/mq75client -g /G/data04/mqm
```

The command installs the OSS components into new `opt/mqm` and `var/mqm` directories in `~install/mq75client`. It installs the Guardian components into `/G/data04/mqm`.

What to do next

For instructions on how to verify your installation, see “Verifying a client installation” on page 374.

Installing an IBM MQ client on HP-UX systems

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

This topic describes the installation of a client, using the **swinstall** program to select which components you want to install. The components and are listed in “Choosing what to install” on page 194 ; you must install at least the Runtime and Client components.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Make your current directory the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Accept the license by running the `mqlicense` script:

```
./mqlicense.sh
```

The license is displayed. If you accept the license, you can continue the installation.

4. Type the following command to start the interactive installation procedure:

```
swinstall -s installation_file
```

installation_file is the absolute path to the installation file. The path must begin with a / and end with the name of the installation file. The installation file has the extension `.v11`.

If the files on your DVD are in uppercase with a “;1” suffix, use this name for the depot.

5. In the resulting menu screen, select **MQSERIES**.
 - a. If you do not want to install all IBM MQ components, open **MQSERIES**
 - 1) Mark the components you want to install. The installer resolves dependencies automatically.
 - 2) Review the information displayed by the installer.
6. Optional: To install IBM MQ to a non-default location, select **Actions > Change Product Location**. For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

7. Select **Actions > Install**. The log file tells you if there are any problems that need fixing.
8. Fix any problems, and click **OK** to install. You are informed when the installation has finished.
9. If this installation is not the first installation on the system, you must enter the following command:

```
swconfig -x allow_multiple_versions=true MQSERIES,1= MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* is the path where you have just installed IBM MQ. If you do not enter this command, the **swlist** command reports the installation as installed instead of configured. You must not use IBM MQ unless the installation is configured.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see *Changing the primary installation*.

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see *setmqenv* and *crtmqenv*.
- For instructions on how to verify your installation, see “Verifying a client installation” on page 374.

Related tasks:

“Uninstalling IBM MQ on HP-UX” on page 391

On HP-UX, you can uninstall the IBM MQ server or client using the **swremove** command.

Non-interactive installation of IBM MQ client on HP-UX

You can perform a non-interactive installation of the IBM MQ client using the **swinstall** command. A non-interactive installation is also known as a silent, or unattended installation.

Before you begin

Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.

About this task

This topic describes the non-interactive installation of a client, using the **swinstall** program to select which components you want to install. The components and are listed in “Choosing what to install” on page 194 ; you must install at least the Runtime and client components.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Make your current directory the location of the installation file. The location might be the mount point of the CD, a network location, or a local file system directory.
3. Accept the IBM MQ license agreement without an interactive prompt by entering the following command:

```
./mqlicense.sh -accept
```
4. Install IBM MQ using the **swinstall** command:
 - a. If this installation is not the first installation on the system, you must add `-x allow_multiple_versions=true` to the **swinstall** command.
 - b. Add the names of the components to install as parameters of the **swinstall** command. The installer automatically resolves any dependencies.
 - c. Optional: Identify the installation location by adding `,l=MQ_INSTALLATION_PATH` as a parameter of the **swinstall** command. For each installation, all of the IBM MQ components that you require must be installed in the same location. The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

For example, to install all IBM MQ components, in a non-default location, as the first installation, enter the following command:

```
swinstall -s /installation_file.v11 MQSERIES,l=/opt/customLocation
```

To perform a partial installation, providing a list of components, in the default location, as the second installation, enter the following command:

```
swinstall -s /installation_file.v11  
MQSERIES.MQM-RUNTIME MQSERIES.MQM-BASE MQSERIES.MQM-CL-HPUX -x allow_multiple_versions=true
```

/installation_file.v11 is the absolute path to the installation file. The path must begin with a / and end with the name of the installation file. The installation file has the extension .v11.

5. If this installation is not the first installation on the system, you must enter the following command:

```
swconfig -x allow_multiple_versions=true MQSERIES,1= MQ_INSTALLATION_PATH
```

where *MQ_INSTALLATION_PATH* is the path where you have just installed IBM MQ. If you do not enter this command, the **swlist** command reports the installation as installed instead of configured. You must not use IBM MQ unless the installation is configured.

What to do next

For instructions on how to verify your installation, see “Verifying a client installation” on page 374.

Installing an IBM MQ client on Linux

Installing an IBM MQ client on a 32 bit or 64 bit Linux system.

Before you begin

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- If this installation is not the first installation on the system, you must ensure that you have write access to /var/tmp.

About this task

This task describes the installation of the client, using the RPM Package Manager installer to select which components you want to install. You must install at least the Runtime and Client components. The components are listed in “Choosing what to install” on page 194.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Make your current directory the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Run the `mqlicense.sh` script. If you want to view a text-only version of license, which can be read by a screen-reader, type:

```
./mqlicense.sh -text_only
```

The license is displayed.

If want to accept the license without it being displayed, you can run the `mqlicense.sh` script with the `-accept` option.

```
./mqlicense.sh -accept
```

You must accept the license agreement before you can proceed with the installation.

4. If this installation is not the first installation on the system, you must run **crtmqpkg** to create a unique set of packages to install on the system:
 - a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that will uniquely identify the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a sub-directory of `/var/tmp/mq_rpms`, in which the unique set of packages is created. The packages have the *suffix* value contained within the filename.

5. Install IBM MQ. The minimum components you must install are the MQSeriesRuntime and the MQSeriesClient.

- To install to the default location, /opt/mqm, use the **rpm -ivh** command to install each component that you require.

For example, to install all components to the default location use the following command:

```
rpm -ivh MQSeries*.rpm
```

If you are using Ubuntu, add the **--force-debian** attribute. For example, to install all components to the default location use the following command:

```
rpm --force-debian -ivh MQSeries*.rpm
```

You must include this option to prevent seeing warning messages from the version of RPM for your platform, which indicates that the RPM packages are not intended to be directly installed using RPM.

- To install to a non-default location use the **rpm --prefix** option. For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.

For example, to install the runtime and server components to /opt/customLocation on a 32-bit Linux system:

```
rpm --prefix /opt/customLocation -ivh MQSeriesRuntime-V.R.M-F.i386.rpm  
MQSeriesClient-V.R.M-F.i386.rpm
```

where:

V Represents the version of the product that you are installing

R Represents the release of the product that you are installing

M Represents the modification of the product that you are installing

F Represents the fix pack level of the product that you are installing

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see Changing the primary installation.

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see setmqenv and crtmqenv.
- For instructions on how to verify your installation, see “Verifying a client installation” on page 374

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Checking the availability of RPM on your machine

You must ensure that RPM is installed on your Linux machine before installing IBM MQ.

About this task

Important: The installation procedure uses the same RPM packages that are used by the other RPM based distributions. Technologies that convert these RPM packages into other forms, such as alien to convert RPMs to Debian packages, are not compatible with the IBM MQ RPM packages and must not be used.

Procedure

1. To determine if the correct RPM package is installed on your system use the following command:

```
dpkg-query -W --showformat '${Status}\n' rpm
```

If you receive a response that is of the form:

```
install ok installed
```

RPM is installed on your system and no further action is required.

If you receive a response that is of the form:

```
unknown ok not-installed
```

RPM is not installed on your system and you must install the RPM package before attempting to install IBM MQ, using the command described in step 2 on page 267.

2. Run the following command, using root authority. In the example, you obtain root authority using the **sudo** command:

```
sudo apt-get install rpm
```

Attention: If this command does not complete successfully, consult your system administrator for instructions specific to your system on how to install the RPM package.

What to do next

You are now ready to install IBM MQ.

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Related information:

setmqinst

Changing the primary installation

Installing an IBM MQ client on Linux Ubuntu

 **V 8.0.0.2** or Linux on POWER Systems - Little Endian  **V 8.0.0.2**

You can install an IBM MQ client on a Linux Ubuntu **V8.0.0.2**, or Linux on POWER Systems - Little Endian, system in accordance with the system requirements web page.

Before you begin

See System Requirements for IBM MQ V8.0 for details of the supported software levels.

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- Ensure that RPM is installed on your system, as RPM is not installed by default on this platform. To determine if the correct RPM package is installed on your system, see “Checking the availability of RPM on your machine” on page 327.
- Once RPM is installed on your system, carry out the following procedure, as root:
 1. Create directory `/etc/rpm`
 2. Add file `/macros`, containing the following code, `%_dbpath /var/lib/rpm`, to the `/etc/rpm` directory.

Attention: You should only set up a `/macros` file if you are not already using RPM, as the previous instruction changes the default system-wide RPM database.

About this task

Install the client by using the RPM Package Manager installer to select the components that you want to install. The components and package names are listed in “Choosing what to install” on page 194.

Procedure

1. Open a shell terminal and set your current directory to the location of the installation packages. The location might be the mount point of the client DVD, a network location, or a local file system directory. You must have root authority to run the following commands. You can do so by adding **sudo** before the following commands, or by changing to the root user in the shell with the **su** command.
2. Run the `mqlicense.sh` script. If you want to view a text-only version of the license, which can be read by a screen reader, type the following message:

```
./mqlicense.sh -text_only
```

The license is displayed.

You must accept the license agreement before you can proceed with the installation.

3. If this installation is not the first installation of IBM MQ on the system, you must run the **crtmqpkg** command to create a unique set of packages to install on the system. For the **crtmqpkg** command to run on Linux, you must install the **pax** command and **rpmbuild**, which is located in the `rpm` package.
 - a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that uniquely identifies the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

Note: This command creates a full copy of the installation packages in a temporary directory. By default, the temporary directory is located at `/var/tmp`. You must ensure that the system has enough free space before you run this command. To use a different location, you can set the `TMPDIR` environment variable before you run the **crtmqpkg** command. For example:

```
TMPDIR=/test ./crtmqpkg
```

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a subdirectory of the `/var/tmp/mq_rpms` directory, in which the unique set of packages

is created. The packages have the *suffix* value contained within the file name. In the following example, the suffix of "1" `./crtmqpkg 1` means that there is a subdirectory named `/var/tmp/mq_rpms/1/i386`.

The packages are renamed according to the subdirectory, for example, on Linux for System x (64-bit):

```
From: MQSeriesRuntime-8.0.0-0.x86_64.rpm
To: MQSeriesRuntime_1-8.0.0-0.x86_64.rpm
```

4. Install IBM MQ. At a minimum, you must install the MQSeriesRuntime component.

An additional flag is required when installing on Linux Ubuntu:

- **--force-debian:** You must include this option to prevent warning messages from the version of RPM for your platform, which indicates that the RPM packages are not intended to be directly installed using RPM.

V 8.0.0.2 An additional flag is required when installing on Linux on POWER Systems - Little Endian

- **--ignorearch:** You must include this option to prevent problems with some levels of rpm not recognizing the Linux on POWER Systems - Little Endian architecture

If you are installing a subset of components, you must ensure that any dependencies are first installed, as listed in Table 58 on page 330.

Note the following:

- To install to the default location, `/opt/mqm`, use the rpm **-ivh** command to install each component that you require.

To install the runtime component to the default location on Ubuntu Linux for System x (64-bit), use the following command:

```
rpm -ivh --force-debian MQSeriesRuntime-*.rpm
```

V 8.0.0.2 To install the runtime component to the default location on Ubuntu Linux on POWER Systems - Little Endian , use the following command:

```
rpm -ivh --force-debian --ignorearch MQSeriesRuntime-*.rpm
```

To install all components to the default location on Ubuntu Linux on POWER Systems - Little Endian use the following command:

```
rpm -ivh --force-debian --ignorearch MQSeries*.rpm
```

- To install to a nondefault location, use the rpm **--prefix** option. For each installation, all of the IBM MQ components that you require must be installed in the same location.

The installation path specified must be, either an empty directory, the root of an unused file system, or a path that does not exist.

Attention: The length of the path is limited to 256 bytes and must not contain spaces.

V 8.0.0.2 For example, enter the following installation path to install the runtime component to the `/opt/customLocation` directory on Ubuntu Linux on POWER Systems - Little Endian :

```
rpm --prefix /opt/customLocation -ivh --force-debian --ignorearch
MQSeriesRuntime-*.rpm
```

Table 58 on page 330 lists all available packages on Ubuntu, together with all the associated dependencies.

To install and use the package listed in the *Package Name* column, you must also install the components listed in the *Package Dependencies* column.

Table 58. Package component dependencies

Package Name	Component Function	Package Dependencies
MQSeriesRuntime	Common function for all other components	None
MQSeriesClient	C IBM MQ client libraries	MQSeriesRuntime
MQSeriesJava	Java and JMS IBM MQ APIs	MQSeriesRuntime
MQSeriesJRE	Java Runtime Environment	MQSeriesRuntime
MQSeriesExplorer	<p>IBM MQ Explorer</p> <p>IBM MQ Explorer is only available on Linux for System x (64-bit).</p> <p>Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2.</p>	MQSeriesRuntime MQSeriesJRE
MQSeriesGSKit	<p>IBM Global Security Kit</p> <p>Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2.</p>	MQSeriesRuntime MQSeriesJRE
MQSeriesSDK	Header files and libraries for non-Java APIs	MQSeriesRuntime
MQSeriesMan	UNIX man pages for IBM MQ	MQSeriesRuntime
MQSeriesSamples	IBM MQ application samples	MQSeriesRuntime
MQSeriesMsg_cs	Language specific message catalog files	MQSeriesRuntime
MQSeriesMsg_de		
MQSeriesMsg_es		
MQSeriesMsg_fr		
MQSeriesMsg_hu		
MQSeriesMsg_it		
MQSeriesMsg_ja		
MQSeriesMsg_ko		
MQSeriesMsg_pl		
MQSeriesMsg_pt		
MQSeriesMsg_ru		
MQSeriesMsg_Zh_CN		
MQSeriesMsg_Zh_TW		

Table 58. Package component dependencies (continued)

Package Name	Component Function	Package Dependencies
MQSeriesFTBase	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesJava MQSeriesJRE
MQSeriesFTLogger	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesFTBase MQSeriesJava MQSeriesJRE
MQSeriesFTTools MQSeriesFTAgent	IBM MQ Managed File Transfer components	MQSeriesRuntime MQSeriesFTBase MQSeriesJava MQSeriesJRE
MQSeriesFTService	IBM MQ Managed File Transfer component	MQSeriesRuntime MQSeriesFTAgent MQSeriesFTBase MQSeriesJava MQSeriesJRE
MQSeriesAMS	Advanced Message Security component Note: There is no IBM support for this component on Ubuntu V 8.0.0.2 , unless you are running on Ubuntu Version 14.04 (or later) and have installed IBM MQ Version 8.0.0.2.	MQSeriesRuntime

Results

You have installed the packages you require.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see “Verifying a client installation” on page 374

Related concepts:

“Multiple installations” on page 183

On UNIX, Linux, and Windows, it is possible to have more than one copy of IBM MQ on a system.

“Choosing a primary installation” on page 184

On systems that support multiple installations of IBM MQ (UNIX, Linux, and Windows), the primary installation is the one to which IBM MQ system-wide locations refer. Having a primary installation is optional, but convenient.

Related tasks:

“Uninstalling IBM MQ on Linux” on page 392

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

Related information:

setmqinst

Changing the primary installation

Installing an IBM MQ client on Solaris

Before you begin

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in “Preparing the system” on page 233.
- This procedure is for the installation of a standard IBM MQ client, from the client DVD. If you are using a server DVD, follow the steps in “Installing IBM MQ server on Solaris” on page 272, and select the appropriate client components in step 8.

About this task

This task describes the installation of the IBM MQ for Solaris server, using the **pkgadd** program. You can choose which components you want to install. The components (or file sets) are listed in “Choosing what to install” on page 194 ; you must install at least the Client component.

Note: If you are installing on the Solaris 11 operating system, ensure that the IPS package (package/svr4) that supports **pkgadd** and equivalent utilities is installed.

Procedure

1. Log in as root, or switch to the superuser using the **su** command.
2. Make your current directory the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Run the `mqlicense.sh` script to accept the license:

```
./mqlicense.sh
```

If you want to view a text-only version of the license, which can be read by a screen-reader, type:

```
./mqlicense.sh -text_only
```

The license is displayed. Follow the instructions to accept the license. If you accept the license, the installation continues. If you do not accept the license, you cannot continue the installation process.

4. If this installation is not the first installation on the system, you must run **crtmqpkg** to create a unique set of packages to install on the system:

- a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that will uniquely identify the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a sub-directory of `/var/spool`, in which the unique set of packages is created. The packages have the *suffix* value contained within the filename.
5. Start the installation process:
 - If the installation is the first installation on the system, enter the following command to start the installation process:

```
pkgadd -d.
```

where " ." means use the current directory.
 - If the installation is not the first installation on the system, enter the following command to start the installation process:

```
pkgadd mqm-suffix
```

where *suffix* is the suffix chosen in the previous step.

6. You are presented with a list of the packages that are available. Enter the number of the `mqm` package.
7. You are prompted to choose a location for installation.
 - To install to the default location, enter `y`.
 - To install to a non-default directory, enter `n`. Then enter the required installation path, and confirm your choice.
8. You receive a number of messages, after which a list of components is displayed. Enter the numbers of the components that you require separated by spaces or commas.
9. If the path chosen in step 7 does not exist, you are asked if you want to create it. You must enter `y` to proceed.
10. Answer any questions appropriately for your system.
11. A message tells you when installation is complete. Enter `q` to exit the `pkgadd` program.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the `setmqenv` or `crtmqenv` command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see [“Verifying a client installation”](#) on page 374.

Related tasks:

[“Uninstalling IBM MQ on Solaris”](#) on page 394

On Solaris, you can uninstall the IBM MQ server or client using the `pkgrm` command.

Non-interactive installation of IBM MQ client on Solaris

Before you begin

- Before you start the installation procedure, make sure that you have completed the necessary steps outlined in [“Preparing the system”](#) on page 233.
- This procedure is for the installation of a standard IBM MQ client, from the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
- If you are using a server DVD, follow the steps in [“Installing IBM MQ server on Solaris”](#) on page 272, and select the appropriate client components in step 8.

About this task

You can perform a silent installation of IBM MQ. A sample script file called `silent.sh` is supplied in the `silent` directory on the DVD. You can use this script to perform a non-interactive installation that requires no input and shows nothing on the screen. It must be run as root.

The installation script `silent.sh` uses an `admin` file and a response file, both of which are supplied in the `silent` directory. You can use these files as supplied to perform a silent installation of all the components, including all the national language features, to the default location.

Note: If you are installing on the Solaris 11 operating system, ensure that the IPS package (package/svr4) that supports **pkgadd** and equivalent utilities is installed.

Procedure

1. Copy the `silent.sh` script into a writeable directory.
2. If this installation is not the first installation on the system, you must run **crtmqpkg** to create a unique set of packages to install on the system:

- a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that will uniquely identify the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a sub-directory of `/var/spool`, in which the unique set of packages is created. The packages have the *suffix* value contained within the filename.

Once a new package has been generated for the second installation the `silent.sh` script needs to have its `MQ_PACKAGE_NAME` variable modified so that its value is not `mqm` but the new package name.

Also the `MQ_PACKAGE_LOCATION` variable needs to be modified so that its value is not `$MQ_MEDIA_LOCATION` but the location of the new package (which by default is `/var/spool/pkg`).

3. Optional: If you want to change where the IBM MQ client DVD is mounted, you must update the values in the `silent.sh` script. By default, the script assumes that the DVD has been mounted at `/CD7FVML`.
4. Optional: If you want to change where the output and logs are written to, update the values in the `silent.sh` script. By default, output and logs are written to the file `/var/tmp/mq.install`.
5. Optional: If you want to install to a non-default location, you must update the `MQ_INSTALLATION_PATH` variable in the `silent.sh` script.

Note:

- The installation path specified must either be an empty directory, the root of an unused file system, or a path that does not exist. The length of the path is limited to 256 bytes and must not contain spaces.
 - If the directory you specified does not exist, the install script creates that directory.
6. Optional: If you want to change the components that are installed, you must edit the response file. A list of all the installable IBM MQ components can be found at: "Choosing what to install" on page 194.

Solaris does not check, during a silent installation, that prerequisite components are installed. You can use the following procedure to create a response file interactively, before using it to install the product. **pkgask** prompts you for the names of the components to install.

- a. Run the **mqlicense.sh** command to accept the license agreement for the product.
- b. **pkgask -d path_to_install_image -r response_file mqm**

The inputs to **pkgask** are the same as those inputs documented for **pkgadd**, but instead of the product being installed a response file is created.

7. Optional: If you have edited the response file, you must then edit the `silent.sh` to use your custom response file.
8. To start the installation, run `silent.sh`.
9. Check the log file for any errors.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see “[Verifying a client installation](#)” on page 374.

Installing an IBM MQ client on Windows systems

This topic describes how to install IBM MQ client on Windows systems. This procedure can be used for installing a first or a subsequent installation.

Before you begin

To install an IBM MQ client, you must be logged on to Windows as an administrator.

About this task

Follow these instructions to perform an interactive compact, typical, or custom installation of IBM MQ. To view all the installation types and the features that are installed with each option consult Features installed with each type of interactive installation.

Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate `setup.exe` in the Windows directory of the IBM MQ installation image.
 - From a DVD, this location might be:
`E:\Windows\setup.exe`
 - From a network location, this location might be:
`m:\instmq\Windows\setup.exe`
 - From a local file system directory, this location might be:
`C:\instmq\Windows\setup.exe`
3. Double-click the **Setup** icon to start the installation process. It is possible to run either by:
 - Running `setup.exe` from the command prompt. Or
 - Double-clicking `setup.exe` from Windows Explorer.

If you are installing on a Windows system with UAC enabled, accept the Windows prompt to allow the launchpad to run as elevated. During installation, you might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the installation to continue.

The IBM MQ Installation window is displayed.

4. Continue to follow the instructions as shown on screen.

Results

A new sample IBM MQ MQI client configuration file is created in the IBM MQ installation directory (for example `C:\Program Files\IBM\WebSphere MQ\`, by the IBM MQ MQI client package, during installation, but only if this file does not exist. This file contains the `ClientExitPath` stanza. An example `mqclient.ini` file is shown in Configuring a client using a configuration file.

Note:

If you are using a common configuration file for multiple clients, either in the IBM MQ installation directory or in another location using the `MQCLNTCF` environment variable, you must grant read access to all user identifiers under which the IBM MQ client applications run. If the file cannot be read, the failure is traced and the search logic continues as if the file had not existed.

What to do next

- If you have chosen this installation to be the primary installation on the system, you must now set it as the primary installation. Enter the following command at the command prompt:

```
MQ_INSTALLATION_PATH\bin\setmqinst -i -p MQ_INSTALLATION_PATH
```

You can have only one primary installation on a system. If there is already a primary installation on the system, you must unset it before you can set another installation as the primary installation. For more information, see [Changing the primary installation](#).

- You might want to set up the environment to work with this installation. You can use the **setmqenv** or **crtmqenv** command to set various environment variables for a particular installation of IBM MQ. For more information, see [setmqenv](#) and [crtmqenv](#).
- For instructions on how to verify your installation, see “[Verifying a client installation](#)” on page 374.

Related concepts:

“[Modifying the client installation on Windows using Add/Remove Programs](#)” on page 347

On some versions of Windows, you can modify an installation by using Add/Remove Programs.

Related tasks:

“[Advanced installation using msiexec](#)”

“[Using the MQParms command](#)” on page 343

“[Uninstalling IBM MQ on Windows systems](#)” on page 395

You can uninstall the IBM MQ MQI clients and servers on Windows systems by using the control panel, the command line (**msiexec**), **MQParms**, or by using the installation media, in which case you can optionally remove queue managers as well.

Advanced installation using msiexec

About this task

IBM MQ on Windows uses the MSI technology to install software. MSI provides both an interactive installation and a non interactive installation. An interactive installation displays panels and ask questions.

The **msiexec** command uses parameters to give MSI some or all of the information that can also be specified through panels during an interactive installation. This means that a user can create a reusable automated or semi-automated installation configuration. Parameters can be given through the command line, a transform file, a response file, or a combination of the three.

Procedure

To install using **msiexec**, at the command line, enter the **msiexec** command in the following format:

```
msiexec parameters [USEINI="response-file"] [TRANSFORMS="transform_file "]
```

Where:

parameters

are either command-line parameters preceded by a / character, or property=value pairs (if using both forms of parameter always put the command-line parameters first). For further information, see “[Specifying command line parameters with msiexec](#)” on page 339.

For an unattended installation, you must include the /q or /qn parameter in the command line. Without this parameter, the installation is interactive.

Note: You must include the `/i` parameter and the file location of the IBM MQ installer package.

response-file

is the full path and file name of the file that contains the [Response] stanza and the required property=value pairs, for example `C:\MyResponseFile.ini`. An example response file, `Response.ini`, is supplied with IBM MQ. This file contains default installation parameters. For further information, see “Using a response file with `msiexec`” on page 340.

transform_file

is the full path and file name of a transform file. For further information, see “Using transforms with `msiexec`” on page 342 and “Multiple installation using MSI Instance ID” on page 285.

Note: For a silent installation to succeed, the `AGREETOLICENSE=?YES?` property must be defined either on the command line or in the response file.

Results

After the command has been entered, the command prompt immediately reappears. IBM MQ is installing as a background process. If you have entered parameters to produce a log, check this file to see how the installation is progressing. If the installation completes successfully, you see the message `Installation operation completed successfully` in the log file.

Specifying command line parameters with `msiexec`

About this task

The `msiexec` command can accept two types of parameters on the command line, as follows:

- Standard command line parameters, preceded by a `/` character.
For a table of the `msiexec` command line parameters, see the MSDN Command-Line Options web page.
- Property=value pair parameters on the command line. All the parameters available for use in a response file can be used on the command line, for a list of these, see Table 60 on page 341. In addition there are some extra property=value pair parameters that are only for use on the command line, for details see Table 59 on page 340.

When using the property=value pair parameters note that:

- Property strings must be in uppercase.
- Value strings are not case-sensitive, except for feature names. You can enclose value strings in double quotation marks. If a value string includes a blank, enclose the blank value string in double quotation marks.
- For a property that can take more than one value, use the format:
`ADDLOCAL="Server,Client"`
- For properties taking paths and filenames, for example `PGMFOLDER`, you must supply the paths as absolute paths and not relative; that is, `C:\folder\file` and not `.\folder\file`.

When using property=value pair and command line parameters with the `msiexec` command, enter command line parameters first.

If a parameter is specified both on the command line and in a response file, the setting on the command line takes precedence.

Example

A typical example of an `msiexec` command is:

```
msiexec /i "path\MSI\IBM WebSphere MQ.msi" /l*v c:\install.log  
/q TRANSFORMS="1033.mst" AGREETOLICENSE="yes" ADDLOCAL="Client"
```

A typical example of an `msiexec` command when you are installing a second copy of the IBM MQ product is:

```
msiexec /i "path\MSI\IBM WebSphere MQ.msi" /! *v c:\install.log
/q TRANSFORMS=":InstanceId2.mst;1033.mst" AGREETOLICENSE="yes"
ADDLOCAL="Client" MSINewInstance=1
```

The following table shows the parameters which can only be provided on the command line and not in a response file.

Table 59. msiexec property=value parameters

Property	Values	Meaning
USEINI	<i>path \ file_name</i>	Use the specified response file. See “Using a response file with msiexec”
SAVEINI	<i>path \ file_name</i>	Generate a response file during installation. The file contains those parameters selected for this installation that a user might make during an interactive installation.
ONLYINI	1 yes ""	1, yes or any value other than null. End the installation before updating the target system, but after generating a response file, if this is specified. "". Continue the installation and update the target system (the default).
TRANSFORMS	:InstanceId <i>x.mst</i> <i>path \ file_name</i> :InstanceId <i>x.mst</i> ; <i>path \ file_name</i>	The :InstanceId <i>x.mst</i> value is only required for a subsequent installation of IBM MQ Version 7.1 or greater. The <i>path \ file_name</i> specifies what transform (.mst) files must be applied to the product. For example, "1033.mst" specifies the supplied U.S. English transform file.
MSINewInstance	1	This property is only required for subsequent installations of IBM MQ Version 7.1 or greater.
REMOVEFEATURES	yes	Required with value "yes" for a silent installation, otherwise ignored. Allows obsolete features, no longer part of IBM MQ, to be deleted.

Using a response file with msiexec

About this task

You can use the **msiexec** command with a parameter which specifies additional properties are defined in a response file. You can combine the msiexec command-line parameters described in “Specifying command line parameters with msiexec” on page 339.

A response file is an ASCII text file, with a format like a Windows .ini file, that contains the stanza [Response]. The [Response] stanza contains some or all the parameters that would normally be specified as part of an interactive installation. The parameters are given in a property=value pair format. Any other stanzas in the response file are ignored by **msiexec**. An example response file, Response.ini, is supplied with IBM MQ. It contains the default installation parameters.

Procedure

A typical example of an msiexec command is: msiexec /i "path\MSI\IBM WebSphere MQ.msi" /! *v c:\install.log TRANSFORMS= "1033.mst" USEINI= "C:\MQ\Responsefile"

If a parameter is specified both on the command line and in a response file, the setting on the command line takes precedence. All the parameters available for use in a response file can also be used on the command line, for a list of these see Table 60.

In the response file, all text is in English, and comments begin with a ; character.

For information about creating a response file, see “Creating a response file” on page 287.

Example

An example of a typical response file:

```
[Response]
PGMFOLDER="c:\mqm"
DATFOLDER="c:\mqm\data"
AGREETOLICENSE="yes"
ADDLOCAL="Client"
REMOVE="Toolkit"
```

Table 60. Response file parameters

Property	Values	Meaning
PGMFOLDER	<i>path</i>	Folder for the IBM MQ program files. For example, c:\mqm.
DATFOLDER	<i>path</i>	Folder for the IBM MQ data files. For example, c:\mqm\data.
USERCHOICE	0 no	<p>If the command line or response file specifies parameters to install features, a dialog can be displayed to prompt you to accept the preselected options, or review and possibly change them.</p> <p>0 or no. Suppresses display of the dialog.</p> <p>Anything else. Dialog is displayed and you can amend the options.</p> <p>Not used for a silent installation.</p>
AGREETOLICENSE	yes	<p>Accept the terms of the license. Set to yes before a silent installation.</p> <p>If the installation is not silent, this parameter is ignored.</p>
ADDLOCAL	<i>feature, feature, All ""</i>	<p>A comma-separated list of features to install locally. For a list of valid feature names, see “IBM MQ features for Windows systems” on page 207.</p> <p>All installs all features</p> <p>"" installs the typical features. If you do not want a feature use REMOVE=" <i>feature</i> "</p> <p>Note: If this is a new installation the typical features (Client, Java, .NET Messaging, and Development Toolkit) are installed by default irrespective of the feature list provided in the ADDLOCAL property. If you do not want a feature use REMOVE="feature"</p>
REMOVE	<i>feature, feature, All ""</i>	<p>A comma-separated list of features to remove. For a list of valid feature names, see “IBM MQ features for Windows systems” on page 207.</p> <p>All uninstalls all features</p> <p>"" uninstalls no features (the default).</p>

Table 60. Response file parameters (continued)

Property	Values	Meaning
INSTALLATIONDESC	?Description of installation?	Sets the installation description from the command line. Subject to the documented installation description length limitations
INSTALLATIONNAME	[INSTALLATION0,]?Name?	Sets the installation name from the command line. Subject to the documented installation name character and length limitations. Note: Supply INSTALLATION0,Name only when upgrading from pre- IBM MQ Version 7.1.
MAKEPRIMARY	0 1 ""	Makes the installation primary, if possible, or removes the primary flag. 1 = Make primary, 0 = Make non-primary, - use default algorithm Note: This option is ignored if a pre- Version 7.1 IBM MQ is installed, or if another Version 7.1 or greater installation is present and set as the primary.

Related tasks:

“Using the MQParms command” on page 343

Related reference:

“Using transforms with msiexec”

Multiple installation using MSI Instance ID

This topic describes how to choose the MSI instance ID you require for non-interactive multiple installations on a client.

About this task

The installation procedure on a client is identical to that on a server. See “Multiple installation using MSI Instance ID” on page 285 for details.

Using transforms with msiexec

MSI can use transforms to modify an installation. During IBM MQ installation, transforms can be used to support different national languages. IBM MQ is supplied with transform files in the \MSI folder of the client image. These files are also embedded in the IBM MQ Windows installer package, IBM WebSphere MQ.msi.

On the **msiexec** command line, you can specify the required language by using the TRANSFORMS property in a property=value pair. For example:

```
TRANSFORMS="1033.mst"
```

You can also specify the full path and file name of the transform file. Again, the quotation marks surrounding the value are optional. For example:

```
TRANSFORMS="D:\Msi\1033.mst"
```

Table 61 on page 343 shows the locale identifier, language, and the transform file name to use in the **msiexec** command line.

You might need to merge transforms to install multiple installations of the same version, for example:

```
TRANSFORMS=":InstanceId2.mst;D:\Msi\1033.mst"
```

You can also specify the required language by using the MQLANGUAGE property with the **MQParms** command. For information about the `msiexec` property=value parameters, see “MQParms parameter file” on page 345.

Parameters

Table 61. Supplied transform files for various language support. This table shows the supplied transform files, the resulting language, and the numeric value to use in the `msiexec` command line.

Language	Transform File name	Value
U.S. English	1033.mst	1033
German	1031.mst	1031
French	1036.mst	1036
Spanish	1034.mst	1034
Italian	1040.mst	1040
Brazilian Portuguese	1046.mst	1046
Japanese	1041.mst	1041
Korean	1042.mst	1042
Simplified Chinese	2052.mst	2052
Traditional Chinese	1028.mst	1028
Czech	1029.mst	1029
Russian	1049.mst	1049
Hungarian	1038.mst	1038
Polish	1045.mst	1045

Creating a response file

A response file is used with `msiexec` on a client.

About this task

The procedure for creating a response file on a client is identical to that on the server. See “Creating a response file” on page 287 for details.

Using the MQParms command

Before you begin

You can use the MQParms command to invoke installation or uninstallation. This command can use parameters on a command line, or those specified in a parameter file. The parameter file is an ASCII text file that contains the parameter values that you want to set for the installation. The MQParms command takes the specified parameters and generates the corresponding `msiexec` command line.

This means that you can save all the parameters that you want to use with the `msiexec` command in a single file.

If you are running IBM MQ on Windows systems with User Account Control (UAC) enabled, you must invoke the installation with elevated privileges. If you are using the Command prompt or MQ Explorer elevate privileges by using a right-click to start the program and selecting **Run as administrator**. If you try to run the MQParms program without using elevated privileges, the installation fails with an error of AMQ4353 in the installation log.

For silent operations, this must include the `/q` or `/qn` parameter, either on the command line, or in the [MSI] stanza of the parameter file. You must also set the `AGREETOLICENSE` parameter to "yes".

You can specify many more parameters in the parameter file that you use with the `MQParms` command than you can in the response file that you use directly with the `msiexec` command. Also, as well as parameters that the IBM MQ installation uses, you can specify parameters that can be used by the Prepare IBM MQ wizard.

If you do not complete the Prepare WebSphere MQ Wizard directly after IBM MQ installations or if for any reason your machine is rebooted between completing IBM MQ installation and completing the Prepare WebSphere MQ Wizard, ensure that the wizard is run with Administrator privilege afterward, otherwise the installation is incomplete, and might fail. You might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the wizard to continue

An example of the file `MQParms.ini` is supplied with IBM MQ. This file contains default installation parameters.

There are two ways to create a parameter file for installation:

- Copy and edit the file `MQParms.ini` that is supplied with the product, using an ASCII file editor.
- Create your own parameter file using an ASCII file editor.

About this task

To invoke installation using the `MQParms` command:

Procedure

1. From a command line, change to the root folder of the IBM MQ client CD (that is, the location of the file `MQParms.exe`).
2. Enter the following command:

```
MQParms [ parameter_file ] [ parameters ]
```

where:

parameter_file

is the file that contains the required parameter values. If this file is not in the same folder as `MQParms.exe`, specify the full path and file name. If you do not specify a parameter file, the default is `MQParms.ini`. For further details, see "MQParms parameter file" on page 345.

parameters

are one or more command-line parameters, for a list of these, see the MSDN Command-Line Options web page.

Example

A typical example of an `MQParms` command is:

```
MQParms "c:\MyParamsFile.ini" /!*v c:\install.log
```

If you specify a parameter both on the command line and in the parameter file, the setting on the command line takes precedence.

If you do not specify `/i`, `/x`, `/a`, or `/j`, `MQParms` defaults to standard installation using the IBM MQ Windows Installer package, `IBM MQ.msi`. That is, it generates the following part of the command line:

```
/i " current_folder \MSI\IBM WebSphere MQ.msi"
```


MQParms parameter file

A parameter file is an ASCII text file that contains sections (stanzas) with parameters that can be used by the MQParms command. Typically, this is an initialization file such as MQParms.ini.

The MQParms command takes parameters from the following stanzas in the file:

[MSI] Contains general properties related to how the MQParms command runs and to the installation of IBM MQ.

The properties that you can set in this stanza are listed in “Advanced installation using msixec” on page 338, and Table 62.

MQParms ignores any other stanzas in the file.

The stanza parameters are in the form `property=value`, where `property` is always interpreted as uppercase, but `value` is case sensitive. If a value string includes a blank, it must be enclosed in double quotation marks. Most other values can be enclosed in double quotation marks. Some properties can take more than one value, for example:

```
ADDLOCAL="Server,Client"
```

To clear a property, set its value to an empty string, for example:

```
REINSTALL=""
```

The following tables show the properties that you can set. The default is shown in bold.

For the [MSI] stanza, you can enter standard MSI command line options and properties. For example:

- /q
- ADDLOCAL="client"
- REBOOT=Suppress

Refer to Table 62, and Table 63 on page 346 for the properties used to install IBM MQ.

Table 62 shows additional properties in the stanza that affect how the MQParms command runs, but that do not affect the installation.

Table 62. Properties used by MQParms in the MSI stanza

Property	Values	Description
MQPLOG	<i>path</i> <i>file_name</i>	MQParms generates a text log file with the specified name and location.
MQPLANGUAGE	system <i>user</i> <i>transform_value</i> <i>existing</i>	The installation language. system . Install using the language of the default system locale (the default). <i>user</i> . Install using the language of the default locale of the user. <i>transform_value</i> . Install using the language specified by this value. See Table 63 on page 346. <i>existing</i> . If MQ already exists on the system, the same language will be used by default, otherwise system is used.
MQPSMS	0 <i>no</i>	0 or <i>no</i> . MQParms does not wait for the <code>msiexec</code> command to end (the default). Any other value. MQParms waits for the <code>msiexec</code> command to end.

Table 62. Properties used by MQParams in the MSI stanza (continued)

Property	Values	Description
MQPINUSE	0 1	If MQPINUSE is set to 1, MQParams continues installing even if IBM MQ files are in use. If this option is used a reboot will be required to complete the installation.

Table 63. Valid values for the MQPLANGUAGE property

Language	Valid values		
U.S. English	English	en_us	1033
German	German	de_de	1031
French	French	fr_fr	1036
Spanish	Spanish	es_es	1034
Italian	Italian	it_it	1040
Brazilian Portuguese		pt_br	1046
Japanese	Japanese	ja_jp	1041
Korean	Korean	ko_kr	1042
Simplified Chinese		zh_cn	2052
Traditional Chinese		zh_tw	1028
Czech	Czech	cs_cz	1029
Russian	Russian	ru_ru	1049
Hungarian	Hungarian	hu_hu	1038
Polish	Polish	pl_pl	1045

A typical example of a parameter file is:

```
[MSI]
MQPLANGUAGE=1033
MQPLOG=%temp%\MQParams.log
MQPSMS=no
ADDLOCAL=CLIENT
/m miffile
REMOVE=""
/l*v c:\install.log
```

Modifying the client installation on Windows

You modify the installation when an IBM MQ for Windows client is installed and you want to remove or install some IBM MQ client features.

1. Insert the IBM MQ client DVD into the DVD drive.
2. If autorun is installed, the installation process starts.
Otherwise, double-click **Setup** in the root folder of the DVD to start the installation process.
The IBM MQ client Setup window is displayed. Click **Next** to continue.

3. Select **Modify**, then click **Next**.

The Features panel is displayed.

4. To change the installation of a feature:
 - a. Click the symbol next to the feature name to display a menu.
 - b. Select the required option from:
 - Install this feature
 - Install this feature and all its subfeatures (if any)
 - Do not install this feature (remove if already installed).

The symbol next to the feature name changes to show the current installation option.

5. When your selections are complete, click **Next**.
6. The IBM MQ client Setup window displays a summary of the installation you selected.
To continue, click **Modify**.
7. Wait until the progress bar is complete.
When the IBM MQ client is successfully installed, the IBM MQ client Setup window displays the following message: Installation Wizard Completed Successfully
Click **Finish** to close the window.

Modifying the client installation on Windows using Add/Remove Programs

On some versions of Windows, you can modify an installation by using Add/Remove Programs.

For Windows 7 follow these steps.

1. From the Windows taskbar, select **Start > Control Panel**.
2. Select **Add/Remove Programs**.
3. Select **IBM MQ**.
4. Select **Change**.

The IBM MQ Setup window with the Program Maintenance panel is displayed. Follow the procedure for modifying the installation by using the process from step 3 to the end.

For Windows 8, the **Add/Remove Programs** option uninstalls the whole product.

You need to run the setup.exe file from the original installation media to make any modifications to the installation.

Silently modifying an IBM MQ client installation using msiexec

To silently modify an installation using msiexec, follow the instructions on the installation pages, but set the ADDLOCAL parameter to include the features you want to add, and set the REMOVE parameter to the features you want to remove.

For example if you used ADDLOCAL="JavaMsg" and REMOVE="" it would modify the installation to include the Java Messaging and Web Services feature.

The instructions for msiexec begin here: "Advanced installation using msiexec" on page 338

Silently modifying an IBM MQ client installation using MQParms

To silently modify an installation using MQParms, follow the instructions on the installation pages, but set the ADDLOCAL parameter to include the features you want to add, and set the REMOVE parameter to the features you want to remove.

For example if you used ADDLOCAL="JavaMsg" and REMOVE="" it would modify the installation to include the Java Messaging and Web Services feature.

For details of the MQParms command, see "Using the MQParms command" on page 287.

Installing an IBM MQ client on IBM i



The IBM MQ client for IBM i is a part of the IBM MQ product.

Before you begin

Note: If you have already installed the IBM MQ Version 8.0 server, you already have a client and must not attempt to install the stand-alone client.

You can install only one instance of IBM MQ client for IBM i in each partition of your server.

When you install IBM MQ client for IBM i two user profiles are created:

- QMQM
- QMQMADM

These two objects are central to the correct running of IBM MQ for IBM i. Do not alter or delete them. If you do, IBM cannot guarantee correct behavior of your product. These profiles are retained when the product is deleted.

About this task

This procedure covers the installation of both the client and the client samples. If you do not want to install the client samples, then do not complete the steps specific to the samples.

After following the optional step to pre-agree the license, and then issuing the **RSTLICPGM** command, the installation runs without requiring any interactive input.

Procedure

1. Sign on to the system with a user profile that has *ALLOBJ special authority, for example QSECOFR.
2. Optional: Pre-agree the license terms and conditions. If you do not choose to pre-agree the license, the license agreement is displayed for you to accept. Run the following commands to pre-agree the license terms and conditions:

- a. For the client:

```
CALL PGM ( QSYS/QLPACAGR ) PARM ( '5725A49' 'V8R0M0' '0000' 0 )
```

The parameters of **PARM** are:

5725A49

The product identifier for IBM MQ client for IBM i

V8R0M0

The version, release, and modification level

0000

The option number for the base IBM MQ client for IBM i product

0 Unused error structure

- b. For the client samples:

```
CALL PGM ( QSYS/QLPACAGR ) PARM ( '5725A49' 'V8R0M0' '0001' 0 )
```

The parameters of **PARM** are:

5725A49

The product identifier for IBM MQ client for IBM i

V8R0M0

The version, release, and modification level

0001

The option number for the samples

0 Unused error structure

3. Issue the installation command to run the installation without requiring any interactive input:

- a. Install the client by issuing the following command:

```
RSTLICPGM LICPGM ( 5725A49 ) DEV ( install device ) OPTION ( *BASE ) OUTPUT ( *PRINT )
```

The parameters of RSTLICPGM are:

LICPGM (5725A49)

The product identifier for IBM MQ client for IBM i

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01

OPTION (*BASE)

The level of IBM MQ client for IBM i product installed

OUTPUT (*PRINT)

Whether the spooled output of the job is printed

- b. Install the samples by issuing the following command:

```
RSTLICPGM LICPGM ( 5725A49 ) DEV ( install device ) OPTION ( 1 ) OUTPUT ( *PRINT )
```

The parameters of RSTLICPGM are:

LICPGM (5725A49)

The product identifier for IBM MQ client for IBM i

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01

OPTION (1)

The samples option

OUTPUT (*PRINT)

Whether the spooled output of the job is printed

4. To ensure that the product has loaded correctly, issue the Display Software Resources (**DSPSFWRSC**) command and check that the licensed program 5725A49 is listed. If you have installed the base and the optional samples, you see:

Resource			
ID	Option	Feature	Description
5725A49	*BASE	5050	IBM MQ client for IBM i
5725A49	1	5050	IBM MQ client for IBM i -Samples

5. To see the library and version number of the products installed, press **F11**, while viewing the Display Software Resources screen. The following screen is displayed:

Resource		Feature			
ID	Option	Feature	Type	Library	Release
5725A49	*BASE	5050	*CODE	QMOM	V8R0M0
5725A49	1	5050	*CODE	QMOMSAMP	V8R0M0

6. To check exactly what version you have installed, use the **DSPMQVER** program. For example, CALL PGM(QMQM/DSPMQVER) from the command line or /QSYS.LIB/QMQM.LIB/DSPMQVER.PGM -a in a qshell.

What to do next

If you want to see how the installation went in more detail, perform one or more of the following tasks:

- View the log file using the DSPJOBLOG command.
- View the spoolfile generated from the RSTLICPGM command.

If the installation of IBM MQ client for IBM i failed, see “Installation failures for IBM i” on page 316

Related concepts:

“Uninstalling IBM MQ for IBM i” on page 400

There are two ways of uninstalling IBM MQ for IBM i.

Installation of IBM MQ client and IBM MQ server for IBM i

When you install an IBM MQ server on an IBM i system, the client is also automatically installed.

The installed version of the IBM MQ client for IBM i can be refreshed by using a "slip installation" which replaces an existing installation with a fresh image.

Installing a client over an existing client results in a successful installation.

Installing a client over an existing server results in a failure with a CPDB6A4 error.

Installing a server over an existing client results in a successful upgrade of the client to both server and client capabilities.

Installing IBM MQ Advanced Message Security

Install and uninstall IBM MQ Advanced Message Security component.

Before you begin

IBM MQ Advanced Message Security is a separately installed and licensed component of IBM MQ and is another option on the IBM MQ installer. Make sure that you purchase a license for using IBM MQ Advanced Message Security before the installation.

Additionally, make sure the following IBM MQ components are installed in your environment:

- MQSeriesRuntime
- MQSeriesServer

Related tasks:

“Installing IBM MQ Advanced Message Security on AIX”

You can install IBM MQ Advanced Message Security component on AIX platforms using either system management interface tool (SMIT) or the command line.

“Installing IBM MQ Advanced Message Security on HP-UX” on page 354

You can install IBM MQ Advanced Message Security component on HP-UX platforms.

“Installing IBM MQ Advanced Message Security on Linux” on page 354

You can install IBM MQ Advanced Message Security on Linux platforms.

“Installing IBM MQ Advanced Message Security on Windows” on page 355

Once you purchase the IBM MQ Advanced Message Security license, you can install the component on Windows platforms.

“Installing IBM MQ Advanced Message Security on IBM i” on page 356

You can install IBM MQ Advanced Message Security component on IBM i.

“Installing IBM MQ Advanced Message Security on z/OS” on page 357

You can install IBM MQ Advanced Message Security component on z/OS.

“Uninstalling IBM MQ Advanced Message Security” on page 405

Information provided guides you through the uninstallation process of IBM MQ Advanced Message Security component.

Installing IBM MQ Advanced Message Security on AIX

You can install IBM MQ Advanced Message Security component on AIX platforms using either system management interface tool (SMIT) or the command line.

Installing using SMIT

Procedure

1. Log on as root.
2. Change the directory to the location of the installation packages.
3. Start the system management interface tool (SMIT). The system management menu is displayed.
4. Select the required SMIT window using the following sequence:
 - Software Installation and Maintenance
 - Install and Update Software
 - Install Software
5. Enter the directory location of the installation package.
6. Press F4 to list the software in the **SOFTWARE name** option.

7. Select the `mqm.ams.rte` and press Enter.
8. Accept the default setting for the remaining options and press Enter.

Results

IBM MQ Advanced Message Security has been installed successfully.

Installing using command line

Procedure

1. Log on as root.
2. Set your current directory to the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. Run the following command:

```
installp -a -c -Y -d. mqm.ams.rte
```

Note the period, signifying the current directory, following the `-d` parameter.

Results

IBM MQ Advanced Message Security component has been installed successfully.

Installing IBM MQ Advanced Message Security on HP-UX

You can install IBM MQ Advanced Message Security component on HP-UX platforms.

Procedure

1. Log on as root.
2. Set your current directory to the location of the installation file. The location might be the mount point of the DVD, a network location, or a local file system directory.
3. In the command line, issue the following command:

```
swinstall -s MQSERIES.MQM-AMS
```

Results

IBM MQ Advanced Message Security component has been installed successfully.

Installing IBM MQ Advanced Message Security on Linux

You can install IBM MQ Advanced Message Security on Linux platforms.

Before you begin

Procedure

1. Log on as root.
2. Set your current directory to the location of the installation file. The location might be the mount point of the server CD, a network share, or a local file system directory.
3. If this installation is not the first installation on the system, you must run the `crtmqpkg` command to create a unique set of packages to install on the system. In order for the `crtmqpkg` command to run on Linux, the `pax` command or `rpmbuild` must be installed.

Important: `pax` and `rpmbuild` are not supplied as part of the product. You must obtain these from your Linux distribution supplier.

- a. Enter the following command:

```
./crtmqpkg suffix
```

where *suffix* is a name of your choosing, that uniquely identifies the installation packages on the system. *suffix* is not the same as an installation name, although the names can be identical. *suffix* is limited to 16 characters in the ranges A-Z, a-z, and 0-9.

Note: This command creates a full copy of the installation packages in a subdirectory of /var/tmp. You must ensure that the system has enough space before running the command.

- b. Set your current directory to the location specified when the **crtmqpkg** command completes. This directory is a subdirectory of /var/tmp/mq_rpms, in which the unique set of packages is created. The packages have the *suffix* value contained within the filename. For example, using a suffix of "1":

```
./crtmqpkg 1
```

there is a subdirectory named /var/tmp/mq_rpms/1/i386 and the packages are renamed, for example:

```
From: MQSeriesAMS-V.R.M-F.i386.rpm  
To: MQSeriesAMS_1-V.R.M-F.i386.rpm
```

where:

- V** Represents the version of the product that you are installing
- R** Represents the release of the product that you are installing
- M** Represents the modification of the product that you are installing
- F** Represents the fix pack level of the product that you are installing

4. In the command line, issue the following command:

This example shows a minimum installation:

```
rpm -iv <package_name>
```

where <package_name> is one of the following:

- MQSeriesAMS-V.R.M-F.i386.rpm
- MQSeriesAMS-V.R.M-F.x86_64.rpm
- MQSeriesAMS-V.R.M-F.ppc.rpm
- MQSeriesAMS-V.R.M-F.s390.rpm

Results

IBM MQ Advanced Message Security has been installed successfully.

Installing IBM MQ Advanced Message Security on Windows

Once you purchase the IBM MQ Advanced Message Security license, you can install the component on Windows platforms.

Using the Launchpad

Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate setup.exe in the base directory of the IBM MQ installation image.
 - From a DVD, this location might be:

`E:\ setup.exe`

- From a network location, this location might be:

`m:\instmq\ setup.exe`

- From a local file system directory, this location might be:

`C:\instmq\ setup.exe`

3. Double-click the **Setup** icon to start the installation process. It is possible to start the process by either:
 - Running `setup.exe` from the command prompt.
 - Double-clicking `setup.exe` from IBM MQ Explorer.

Note: If you are installing on a Windows system with UAC enabled, accept the Windows prompt to allow the launchpad to run as elevated. During installation, you might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the installation to continue.

The IBM MQ Installation Launchpad window is displayed.

4. Continue to follow the Launchpad instructions as shown on screen.

Installing IBM MQ Advanced Message Security on IBM i

You can install IBM MQ Advanced Message Security component on IBM i.

Procedure

Install IBM MQ AMS using the command:

```
RSTLICPGM LICPGM(5724H72) DEV(install device) OPTION(2) OUTPUT(*PRINT)
```

where the parameters of **RSTLICPGM** are:

LICPGM(5724H72)

The product identifier for IBM MQ for IBM i.

DEV(install device)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

OPTION(2)

Install IBM MQ Advanced Message Security for IBM i

OUTPUT(*PRINT)

The output is printed with the spooled output of the job.

Results

The IBM MQ AMS component has been installed successfully.

Once IBM MQ AMS is installed on an IBM MQ server installation, any:

- Queue managers that are subsequently started enable security policy management features.
- Applications that connect to the queue manager enable interceptors.

What to do next

See IBM i: Setting up certificates and the keystore configuration file for details on setting up your security policy.

Installing IBM MQ Advanced Message Security on z/OS

You can install IBM MQ Advanced Message Security component on z/OS.

Procedure

IBM MQ Advanced Message Security for z/OS is installed separately using SMP/E by following the process documented in its program directory.

Results


IBM MQ Advanced Message Security component has been installed successfully.

Related concepts:

“IBM MQ Advanced Message Security for z/OS” on page 416

IBM MQ Advanced Message Security for z/OS (AMS) is a separately licensed enabling product that extends IBM MQ to provide a high level of protection for sensitive data flowing through the IBM MQ network using a public key cryptography model.

Related information:

 [IBM MQ Advanced Message Security for z/OS V8.0 Program Directory \(GI13-3329-00\)](#)

Installing IBM MQ Java messaging and web services for IBM i

Install IBM MQ Java messaging and web services for IBM i from either product CD, using the **RSTLICPGM** command.

Before you begin

You can install only one instance of IBM MQ Client for IBM i in each partition of your server.

If you have Java messaging and web services v7.0 or v7.1 installed and want to install v8.0, you can install the new version without uninstalling the old one.

If you have MA88 installed, and try to install anyway, the installation fails with a warning requesting you to uninstall the old client. To uninstall MA88, issue the following command:

```
DLTLICPGM LICPGM(5648C60) OPTION(*ALL)
```

If this command fails to delete the IFS directory /QIBM/ProdData/mqm/java and its subdirectories, use the EDTF command and select option 9 against the Java directory. For example:

```
EDTF STMF('/QIBM/ProdData/mqm')
```

About this task

This procedure covers the installation of both the Java messaging and web services, and the Java messaging and web services samples. If you do not want to install the samples, then do not complete the steps specific to the samples.

After following the optional step to pre-agree the license, and then issuing the **RSTLICPGM** command, the installation runs without requiring any interactive input.

Procedure

1. Sign on to the system with a user profile that has *ALLOBJ special authority, for example QSECOFR.
2. Optional: Pre-agree the license terms and conditions. If you do not choose to pre-agree the license, the license agreement is displayed for you to accept. Run the following commands to pre-agree the license terms and conditions:

- a. For Java messaging and web services:

```
CALL PGM ( QSYS/QLPACAGR ) PARM ( '5724L26' 'V8R0M0' '0000' 0 )
```

The parameters of **PARM** are:

5724L26

The product identifier for IBM MQ Java messaging and web services for IBM i

V8R0M0

The version, release, and modification level

0000

The option number for the base IBM MQ Java messaging and web services product.

0 Unused error structure

- b. For the samples:

```
CALL PGM ( QSYS/QLPACAGR ) PARM ( '5724L26' 'V8R0M0' '0001' 0 )
```

The parameters of **PARM** are:

5724L26

The product identifier for IBM MQ Java messaging and web services for IBM i

V8R0M0

The version, release, and modification level

0001

The option number for the samples.

0 Unused error structure

3. Issue the installation command to run the installation without requiring any interactive input:
 - a. Install the IBM MQ Java messaging and web services by issuing the following command:

```
RSTLICPGM LICPGM ( 5724L26 ) DEV ( install device ) OPTION ( *BASE ) OUTPUT ( *PRINT )
```

The parameters of RSTLICPGM are:

LICPGM (5724L26)

The product identifier for IBM MQ Java messaging and web services for IBM i

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01

OPTION (*BASE)

Install the base IBM MQ Java messaging and web services for IBM i

OUTPUT (*PRINT)

Whether the spooled output of the job is printed

- b. Install the samples by issuing the following command:

```
RSTLICPGM LICPGM ( 5724L26 ) DEV ( install device ) OPTION ( 1 ) OUTPUT ( *PRINT )
```

The parameters of RSTLICPGM are:

LICPGM (5724L26)

The product identifier for IBM MQ Java messaging and web services for IBM i

DEV (*install device*)

The device from which the product is to be loaded, typically an optical drive, for example, OPT01

OPTION (1)

Install the samples

OUTPUT (*PRINT)

Whether the spooled output of the job is printed

4. To ensure that the product has loaded correctly, issue the Display Software Resources (DSPSWRSC) command and check that the licensed program 5724L26 is listed. If you have installed the base and the optional samples, you see:

Resource				
ID	Option	Feature	Description	
5724L26	*BASE	5050	IBM MQ Java Messaging and Web Services	
5724L26	1	5050	IBM MQ Java Messaging and Web Services - Samp	

5. Press **F11** while viewing the Display Software Resources screen, and you see the library and version number of the products installed:

Resource					
ID	Option	Feature	Type	Library	Release
5724L26	*BASE	5050	*CODE	QMOMJAVA	V8R0V0
5724L26	1	5050	*CODE	QMOMJAVA	V8R0V0

6. Check what versions you have installed by using the following commands:

IBM MQ Classes for Java:

```
java com.ibm.mq.MQJavaLevel
```

Note: For this command to work, you might have to set your environment classpath to:

- /QIBM/ProdData/mqm/java/lib/com.ibm.mq.jar

IBM MQ Classes for Java Message Service:

```
java com.ibm.mq.jms.MQJMSLevel
```

Note: For this command to work, you might need to set your environment classpath to:

- /QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar

See Environment variables relevant to IBM MQ classes for Java and Environment variables used by IBM MQ classes for JMS.

For v8.0, both report:

```
Version: 8.0.0.0
```

Note: The command uses the Java classes, and so it reports the version and also performs some verification that the classes are installed and working.

7. See the following topics for full details of verification of both:
 - Using IBM MQ classes for Java
 - Using IBM MQ classes for JMS

Verifying an IBM MQ installation

distributed

The topics in this section provide instructions on how to verify a server or a client installation of IBM MQ on distributed systems.

To verify a server installation, either using the command line or using the postcard application, see “Verifying a server installation.”

To verify a client installation, either using the command line or using the MQ Explorer, see “Verifying a client installation” on page 374.

Related concepts:

“Installing IBM MQ” on page 249

The topics in this section provide instructions on how to install IBM MQ.

“Uninstalling” on page 389

The topics in this section provide instructions on how to uninstall components.

Verifying a server installation

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

You can use either the command line or the postcard application to verify your installation. The postcard application is Java based and requires a system with the ability to view a graphical display.

A local installation uses a single queue manager while a server-to-server installation has multiple queue managers and queues, and both sender and receiver channels.

For a server-to-server verification, the communication links between the two systems must be checked. Before you can do the verification, you must ensure that the communications protocol is installed and configured on both systems. The examples explain how to verify your installation using TCP.

UNIX systems

IBM MQ supports both TCP and SNA. If you do not use TCP, see Setting up communication on UNIX and Linux systems.

Linux IBM MQ for Linux supports TCP on all Linux platforms. On x86 platforms and Power platforms, SNA is also supported. If you want to use the SNA LU6.2 support on these platforms, you need the IBM Communications Server for Linux Version 6.2. The Communications Server is available as a PRPQ product from IBM. For more details, see <http://www.ibm.com/software/network/commserver/about>.

If you do not use TCP, see Setting up communication on UNIX and Linux systems.

Windows

IBM MQ for Windows supports TCP, SNA, NetBios, and SPX. If you do not use TCP, see Setting up communication for Windows .

Related concepts:

“Verify the installation using the command line”

You can use the command line to verify a local installation or a server-to-server installation.

“Verify the installation using the Postcard application” on page 370

You can set up and use the Postcard application to verify a local installation or a server-to-server installation.

Related tasks:

“Verifying a local installation using the command line”

For distributed platforms, you can verify a local installation by using the command line to create a simple configuration of one queue manager and one queue.

“Verifying a server-to-server installation using the command line” on page 367

You can verify a server-to-server installation using two servers, one as a sender and one as a receiver.

“Using the Postcard application to verify a local installation” on page 370

Sending messages successfully between two Postcard applications verifies a local installation.

“Using the Postcard application to verify a server-to-server installation” on page 372

You can use two instances of the Postcard application to verify that a server-to-server installation is working.

Verify the installation using the command line

You can use the command line to verify a local installation or a server-to-server installation.

Use the command line to verify that IBM MQ is successfully installed, and that the associated communication links are working properly.

You can also verify an installation using the postcard application. The postcard application is Java based and requires a system with the ability to view a graphical display. See “Verify the installation using the Postcard application” on page 370.

Related tasks:

“Verifying a local installation using the command line”

For distributed platforms, you can verify a local installation by using the command line to create a simple configuration of one queue manager and one queue.

“Verifying a server-to-server installation using the command line” on page 367

You can verify a server-to-server installation using two servers, one as a sender and one as a receiver.

Verifying a local installation using the command line

For distributed platforms, you can verify a local installation by using the command line to create a simple configuration of one queue manager and one queue.

Before you begin

This procedure is for distributed platforms.  To verify a z/OS installation, see Running the basic installation verification program.

To verify the installation, you must first install the samples package.

Before beginning the verification procedure, you might want to check that you have the latest fixes for your system. For more information about where to find the latest updates, see “Finding product requirements and updated support information” on page 225.

About this task

Use the following steps to configure your default queue manager from the command line. After the queue manager is configured, use the `amqsput` sample program to put a message on the queue. You then use the `amqsget` sample program to get the message back from the queue.

IBM MQ object definitions are case-sensitive. Any text entered as an MQSC command in lowercase is converted automatically to uppercase unless you enclose it in single quotation marks. Make sure that you type the examples exactly as shown.

Procedure

1. If you are verifying an installation on UNIX, Linux or IBM i, log in as a user in the `mqm` group.
2. If you are verifying an installation on UNIX, Linux or Windows, set up your environment:

- a. Set up environment variables for use with a particular installation by entering one of the following commands:

- On Windows:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.

- On UNIX and Linux:

```
. MQ_INSTALLATION_PATH/bin/setmqenv -s
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.

- b. Check that the environment is set up correctly by entering the following command:

```
dspmqver
```

If the command completes successfully, and the expected version number and installation name are returned, the environment is set up correctly.

3. Create a queue manager called QMA by entering the following command:

- On UNIX, Linux, and Windows:

```
crtmqm QMA
```

- On IBM i:

```
CRTMQM MQMNAME(QMA) DFTQMGR(*YES)
```

The **DFTQMGR(*YES)** option specifies that this queue manager is the default queue manager.

Messages indicate when the queue manager is created, and when the default IBM MQ objects are created.

4. Start the queue manager, by entering one of the following commands:

- On UNIX, Linux, and Windows:

```
strmqm QMA
```

- On IBM i:

```
STRMQM MQMNAME(QMA)
```

A message indicates when the queue manager starts.

5. On UNIX, Linux, and Windows, start MQSC by entering the following command:

```
runmqsc QMA
```

A message indicates when MQSC starts. MQSC has no command prompt.

6. Define a local queue called QUEUE1 by entering one of the following commands:

- On UNIX, Linux, and Windows:

```
DEFINE QLOCAL (QUEUE1)
```

- On IBM i:
CRTMQMQ QNAME(Queue1) QTYPE(*LCL)

A message indicates when the queue is created.

7. On UNIX, Linux, and Windows, stop MQSC by entering the following command:
end

Messages are shown, followed by the command prompt.

Note: Subsequent steps require that the samples package is installed.

8. If you are verifying an installation on a UNIX or Linux system, change into the `MQ_INSTALLATION_PATH/samp/bin` directory, which contains the sample programs. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.
9. Put a message on the queue by entering one of the following commands:

- On UNIX and Linux:
./amqsput Queue1 QMA
- On Windows:
amqsput Queue1 QMA
- On IBM i:
CALL PGM(QMQM/AMQSPUT0) PARM(Queue1)

The following messages are shown:

```
Sample AMQSPUT0 start  
target queue is Queue1
```

10. Type some message text on one or more lines, where each line is a different message. Enter a blank line to end the message input. The following message is shown:
Sample AMQSPUT0 end

Your messages are now on the queue and the command prompt is shown.

11. Get the messages from the queue, by entering one of the following commands:
 - On UNIX and Linux:
./amqsget Queue1 QMA
 - On Windows:
amqsget Queue1 QMA
 - On IBM i:
CALL PGM(QMQM/AMQSGET4) PARM(Queue1)

The sample program starts, and your messages are displayed.

Results

You have successfully verified your local installation.

Verifying a server-to-server installation using the command line

You can verify a server-to-server installation using two servers, one as a sender and one as a receiver.

Before you begin

- Make sure that TCP/IP and IBM MQ are installed on both servers.
- Make sure that you are a member of the IBM MQ administrators group (`mqm`) on each server.
- Decide which installation is the sender server and which installation is the receiver server. The installations might be on the same system, or on different systems.

About this task

This procedure provides instructions for Windows, UNIX and Linux systems only.

IBM MQ object definitions are case-sensitive. Any text entered as an MQSC command in lowercase is converted automatically to uppercase unless you enclose it in single quotation marks. Make sure that you type the examples exactly as shown.

Procedure

1. On the **receiver** server:

- a. If the receiver server is a UNIX or Linux system, log in as a user in the `mqm` group.
- b. Check which ports are free, for example by running **netstat**. For more information about this command, see the documentation of your operating system.

If port 1414 is not in use, make a note of 1414 to use as the port number in step 2 h. Use the same number for the port for your listener later in the verification. If it is in use, note a port that is not in use; for example 1415.

- c. Set up the environment for the installation you are using by entering one of the following commands at the command prompt:

- On Windows:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.

- On UNIX and Linux systems:

```
. MQ_INSTALLATION_PATH/bin/setmqenv -s
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.

- d. Create a queue manager called QMB by entering the following command at the command prompt:

```
crtmqm QMB
```

Messages tell you that the queue manager has been created, and that the default IBM MQ objects have been created.

- e. Start the queue manager by entering the following command:

```
strmqm QMB
```

A message tells you when the queue manager has started.

- f. Start MQSC by entering the following command:

```
runmqsc QMB
```

A message tells you that MQSC has started. MQSC has no command prompt.

- g. Define a local queue called RECEIVER.Q by entering the following command:

```
DEFINE QLOCAL (RECEIVER.Q)
```

A message tells you the queue has been created.

- h. Define a listener by entering the following command:

```
DEFINE LISTENER (LISTENER1) TRPTYPE (TCP) CONTROL (QMGR) PORT ( PORT_NUMBER )
```

Where *port_number* is the name of the port the listener runs on. This number must be the same as the number used when defining your sender channel.

- i. Start the listener by entering the following command:

```
START LISTENER (LISTENER1)
```

Note: Do not start the listener in the background from any shell that automatically lowers the priority of background processes.

- j. Define a receiver channel by entering the following command:

```
DEFINE CHANNEL (QMA.QMB) CHLTYPE (RCVR) TRPTYPE (TCP)
```

A message tells you when the channel has been created.

- k. End MQSC by typing:

```
end
```

Some messages are displayed, followed by the command prompt.

2. On the **sender** server:

- a. If the sender server is a UNIX or Linux system, log in as a user in the `mqm` group.
b. Set up the environment for the installation you are using by entering one of the following commands at the command prompt:

- On Windows:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where *MQ_INSTALLATION_PATH* refers to the location where IBM MQ is installed.

- On UNIX and Linux systems:

```
. MQ_INSTALLATION_PATH/bin/setmqenv -s
```

where *MQ_INSTALLATION_PATH* refers to the location where IBM MQ is installed.

- c. Create a queue manager called QMA by entering the following command at the command prompt:

```
crtmqm QMA
```

Messages tell you that the queue manager has been created, and that the default IBM MQ objects have been created.

- d. Start the queue manager, by entering the following command:

```
strmqm QMA
```

A message tells you when the queue manager has started.

- e. Start MQSC by entering the following command:

```
runmqsc QMA
```

A message tells you that an MQSC session has started. MQSC had no command prompt.

- f. Define a local queue called QMB (to be used as a transmission queue) by entering the following command:

```
DEFINE QLOCAL (QMB) USAGE (XMITQ)
```

A message tells you when the queue has been created.

- g. Define a local definition of the remote queue with by entering the following command:


```
DEFINE QREMOTE (LOCAL.DEF.OF.REMOTE.QUEUE) RNAME (RECEIVER.Q)
RQMNAME ('QMB') XMITQ (QMB)
```

- h. Define a sender channel by entering one of the following commands:

con-name is the TCP/IP address of the receiver system. If both installations are on the same system, the *con-name* is localhost. *port* is the port you noted in 1 b. If you do not specify a port, the default value of 1414 is used.

```
DEFINE CHANNEL (QMA.QMB) CHLTYPE (SDR)
CONNNAME ('CON-NAME(PORT)') XMITQ (QMB) TRPTYPE (TCP)
```

- i. Start the sender channel by entering the following command:

```
START CHANNEL(QMA.QMB)
```

The receiver channel on the receiver server starts automatically when the sender channel starts.

- j. Stop MQSC by entering the following command:

```
end
```

Some messages are displayed, followed by the command prompt.

- k. If the sender server is a UNIX or Linux system, change into the *MQ_INSTALLATION_PATH/samp/bin* directory. This directory contains the sample programs. *MQ_INSTALLATION_PATH* represents the high-level directory in which IBM MQ is installed.
- l. If both the sender server and receiver server are installations on the same system, check that the queue managers have been created on different installations by entering the following command:
- ```
dspmqr -o installation
```

If the queue managers are on the same installation, move either QMA to the sender installation or QMB to the receiver installation by using the **setmqm** command. For more information, see **setmqm**.

- m. Put a message on the local definition of the remote queue, which in turn specifies the name of the remote queue. Enter one of the following commands:

- On Windows:  

```
amqsput LOCAL.DEF.OF.REMOTE.QUEUE QMA
```
- On UNIX and Linux systems:  

```
./amqsput LOCAL.DEF.OF.REMOTE.QUEUE QMA
```

A message tells you that **amqsput** has started.

- n. Type some message text on one or more lines, followed by a blank line. A message tells you that **amqsput** has ended. Your message is now on the queue and the command prompt is displayed again.

### 3. On the **receiver** server:

- a. If your receiver server is a UNIX or Linux system, change into the *MQ\_INSTALLATION\_PATH/samp/bin* directory. This directory contains the sample programs. *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.
- b. Get the message from the queue on the receiver by entering one of the following commands:

- On Windows:  

```
amqsget RECEIVER.Q QMB
```
- On UNIX and Linux systems:  

```
./amqsget RECEIVER.Q QMB
```

The sample program starts, and your message is displayed. After a pause, the sample ends. Then the command prompt is displayed.

## Results

You have now successfully verified the server-to-server installation.

## Verify the installation using the Postcard application

You can set up and use the Postcard application to verify a local installation or a server-to-server installation.

Use the Postcard application to verify that IBM MQ is successfully installed, and that the associated communication links are working properly.

The postcard application is Java based and requires a system with the ability to view a graphical display. You can also verify an installation using the command line, see “Verify the installation using the command line” on page 364.

**Note:** Using Postcard to verify an IBM MQ installation is only possible if there is one IBM MQ installation on that box. The Default Configuration wizard will not create a default configuration if a queue manager already exists on the box. The Default Configuration wizard will run on any installation on a box but only one default configuration can be created per box. Using Postcard to verify second and subsequent installations of IBM MQ on the same box is not possible.

### Related tasks:

“Using the Postcard application to verify a local installation”

Sending messages successfully between two Postcard applications verifies a local installation.

“Using the Postcard application to verify a server-to-server installation” on page 372

You can use two instances of the Postcard application to verify that a server-to-server installation is working.

## Using the Postcard application to verify a local installation

Sending messages successfully between two Postcard applications verifies a local installation.

### Before you begin

To verify that the local installation is working, you can run two instances of the Postcard application on the same server. The postcard application can send messages to, and receive messages from, other postcard applications. Successful sending and receiving of messages verifies that IBM MQ is installed and working correctly on the server.

### Note:

- If the system has multiple IBM MQ installations, ensure that Postcard has not been run before on any installations on that server. As the default configuration can only exist on one IBM MQ installation per system, the Default Configuration wizard and Postcard can not be used for verification of a second or any subsequent installation.
- The Postcard application has a graphical interface. To view this interface, your system requires the ability to view a graphical display.
- Before you can run the Postcard application, you must ensure that you are a member of the IBM MQ administrators group ( `mqm` ).

### Procedure

1. If you are verifying an installation on UNIX, Linux or Windows systems, log on as a user in group `mqm`.
2. Start the postcard application in one of the following ways:
  - a. From the command line:
    - 1) Change the directory to `MQ_INSTALLATION_PATH/java/bin`. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.
    - 2) Run the postcard application by entering one of the following commands:
      - UNIX and Linux systems:

- ./postcard
  - Windows systems:
    - postcard
- b. From the IBM MQ Explorer: On Windows and Linux systems (x86-64 platforms), you can start MQ Explorer by using the system menu, the **MQExplorer** command (preferred command), or the MQExplorer executable file. The **strmqcfcg** command is still usable.
    - 1) If the Welcome to IBM MQ Explorer Content view page does not show, click **IBM MQ** in the Navigator view to show the Welcome page.
    - 2) Click **Launch Postcard** to start the Postcard.
  3. At the Postcard - Sign On window, type in a nickname to use to send messages within the Postcard application (for example, User1).
  4. Select the queue manager to use as the mailbox:
    - If you do not have any queue managers, you are prompted to either launch the Default Configuration or close the Postcard application. Launching the Default Configuration creates a default queue manager.
    - If the only queue manager on your server is the default queue manager, this queue manager is used automatically for the postcard application. The default queue manager is created by running the Default Configuration wizard
    - If you have created your own queue managers, but you have not run the Default Configuration wizard, select an appropriate queue manager from the list.
    - If you have run the Default Configuration wizard and you want to use the default queue manager, but there are other queue managers on your server, select the **Advanced** check box. Then select **Use Default Configuration as mailbox**.
    - If you have run the Default Configuration wizard and also created your own queue managers, and you do not want to use the default queue manager, select the **Advanced** check box. Then select **Choose queue manager as mailbox**, and then select the appropriate queue manager from the list.

When your selection is complete, click **OK** to display your first Postcard window.

5. Run a second instance of the Postcard application by following the steps used to open the first instance of the Postcard application.
6. The Postcard - Sign On panel is displayed again. Type in a second nickname to use to send messages within this second Postcard application (for example, User2).
7. Repeat the selection of the queue manager that you want to use as the mailbox (as described in step 4). The queue manager you select for this second Postcard must be the same queue manager as used for the first instance of the Postcard application.
8. In the first Postcard, (User1), enter the nickname ( User2) for the second Postcard application in the **To:** field. Because the sender and receiver are on the same server, you can leave the **On:** field blank.
9. Type a message in the **Message:** field and click **Send**.
10. The **Postcards sent and received** area of the Postcard shows details of the message. In the sending Postcard, the message is displayed as sent. In the receiving Postcard, the message is displayed as received.
11. In the receiving Postcard, (User2), double-click the message in the **Postcards sent and received** area to view it. When this message arrives, it verifies that IBM MQ is correctly installed.

## What to do next

Depending on your situation, you might want to do the following tasks:

- Install IBM MQ on other servers. Follow the installation procedure for the appropriate platform. Ensure that you use the **Join Default Cluster** window in the Default Configuration wizard to add the other servers to the cluster on your first server.

- Install the IBM MQ MQI client on other servers. See “Installing an IBM MQ client” on page 319.
- Continue with further administration tasks, see Administering IBM MQ .

## Using the Postcard application to verify a server-to-server installation

You can use two instances of the Postcard application to verify that a server-to-server installation is working.

### Before you begin

You can use the Postcard application on two servers, one instance of the Postcard application on each server, to verify that a server-to-server installation is working. Successful sending and receiving of messages verifies that IBM MQ is successfully installed, and that communication between the two servers is working correctly.

#### Note:

- If the system has multiple IBM MQ installations, ensure that Postcard has not been run before on any installations on that server. As the default configuration can only exist on one IBM MQ MQ installation per system, the Default Configuration wizard and Postcard can not be used for verification of a second or any subsequent installation.
- The two server installations must be on different systems to do a server-to-server verification using the postcard application. To verify a server-to-server installation on the same machine, you can use the command line. See “Verifying a server-to-server installation using the command line” on page 367
- Make sure that TCP/IP and IBM MQ are installed on both servers.
- Make sure that your systems are able to view a graphical display.
- Make sure that you are a member of the IBM MQ administrators group ( `mqm` ) on each server.
- Check that one of the following scenarios applies:
  - Neither server has had any queue managers created.
    - Use the Default Configuration wizard to create default queue managers on each server and link them to the default cluster.
      - Details on how to use the Default Configuration wizard are provided in this topic.
    - Both servers have existing queue managers and these queue managers are in the same cluster.
      - If your queue managers are not in the same cluster, create new queue managers on both servers. Then create a cluster, and ensure that the queue managers that you create on each server belong to that cluster.
    - You have configured channels to communicate between the two servers.
      - For instructions on how to set up the channels, see “Verifying a server-to-server installation using the command line” on page 367. After you have set up the channels, follow the instructions in this topic to verify your server-to-server installation.

### Procedure

1. On the first server, if you are verifying an installation on UNIX or Linux systems, log on as a user in group `mqm`.
2. Start the postcard application in one of the following ways:
  - a. From the command line:
    - 1) Change the directory to `MQ_INSTALLATION_PATH/java/bin`. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.
    - 2) Run the postcard application by entering one of the following commands:
      - UNIX and Linux systems:
 

```
./postcard
```
      - Windows systems:

postcard


- b. From the IBM MQ Explorer: On Windows and Linux systems (x86-64 platforms), you can start MQ Explorer by using the system menu, the MQExplorer executable file, or the **strmqcfcg** command.
    - 1) If the Welcome to IBM MQ Explorer Content view page does not show, click **IBM MQ** in the Navigator view to show the Welcome page.
    - 2) Click **Launch Postcard** to start the Postcard.
  3. At the Postcard - Sign On window, type a nickname to use to send messages within the Postcard application. For example, User1 for the first server, and User2 for the second server.
  4. When you have completed the wizard, you are taken back to the Postcard - Sign On window.
  5. Select the queue manager to use as the mailbox:
    - If you do not have any queue managers, you are prompted to either launch the Default Configuration or close the Postcard application. Work through the Default Configuration wizard. When you get to the option to join the queue manager to the default cluster, tick the check box. On the next screen:
      - For the first server, select **yes, make it the repository for the cluster**.
      - For the second server, select **No another computer has already joined the cluster as a repository**. When requested, enter the location of the repository, by typing the name of the sender server.
    - If the only queue manager on your server is the default queue manager, this queue manager is used automatically for the postcard application. The default queue manager is created by running the Default Configuration wizard
    - If you have created your own queue managers, but you have not run the Default Configuration wizard, select an appropriate queue manager from the list.
    - If you have run the Default Configuration wizard and you want to use the default queue manager, but there are other queue managers on your server, select the **Advanced** check box. Then select **Use Default Configuration as mailbox**.
    - If you have run the Default Configuration wizard and also created your own queue managers, and you do not want to use the default queue manager, select the **Advanced** check box. Then select **Choose queue manager as mailbox**, and then select the appropriate queue manager from the list.

When your selection is complete, click **OK**.
  6. Complete steps 1 - 5 for the second server.
  7. In the Postcard on the first server:
    - a. Enter the nickname ( user2) for the Postcard application on the second server in the **To:** field.
    - b. Enter the queue manager on the second server in the **On:** field.
    - c. Type a message in the **Message:** field and click **Send**.
  8. In the Postcard on the second server:
    - a. In the **Postcards sent and received**, double-click the message marked as received to view the message from the first server.
    - b. Optional: Send a postcard to the first server by adapting the instructions in step 7. You must enter details of the first server in the **To:** field and the **On:** field.
- The messages verify that IBM MQ is correctly installed and that your communication link between the two servers is working correctly.

---

## Verifying a client installation

You can verify that your IBM MQ MQI client installation completed successfully and that the communication link is working.

The verification procedure shows how to create a queue manager called `queue.manager.1`  (not on z/OS), a local queue called `QUEUE1`, and a server-connection channel called `CHANNEL1` on the server.

It shows how to create the client-connection channel on the IBM MQ MQI client workstation. It then shows how to use the sample programs to put a message onto a queue, and get the message from the queue.

The example does not address any client security issues. See [Setting up IBM MQ MQI client security](#) for details if you are concerned with IBM MQ MQI client security issues.

The verification procedure assumes that:

- The full IBM MQ server product has been installed on a server.
- The server installation is accessible on your network.
- The IBM MQ MQI client software has been installed on a client system.
- The IBM MQ sample programs have been installed.
- TCP/IP has been configured on the server and client systems. For more information, see [Configuring connections between the server and client](#).

To begin the verification procedure using the command line, see [“Verifying a client installation using the command line.”](#)

To begin the verification procedure for the Windows and Linux operating systems when you are using the IBM MQ Explorer, see [“Verifying a client installation using IBM MQ Explorer”](#) on page 379.

### Related tasks:

[“Setting up the server using the command line”](#) on page 375

Follow these instructions to create a queue manager, queue, and channel on the server. You can then use these objects to verify the installation.

[“Connecting an IBM MQ MQI client to a queue manager, using the MQSERVER environment variable”](#) on page 378

When an IBM MQ application is run on the IBM MQ MQI client, it requires the name of the MQI channel, the communication type, and the address of the server to be used. Provide these parameters by defining the `MQSERVER` environment variable.

[“Setting up the server using IBM MQ Explorer”](#) on page 380

You can use the IBM MQ Explorer to create a queue manager, queue and server-connection channel on Windows and Linux systems .

[“Setting up the client using IBM MQ Explorer”](#) on page 381

You can use MQ Explorer to define the client-connection if you are setting up the client and server on the same workstation on a Windows or Linux system.

[“Testing communication between a client and a server”](#) on page 382

On the IBM MQ MQI client workstation, use the `amqsputc` sample program to put a message on the queue at the server workstation. Use the `amqsgetc` sample program to get the message from the queue back to the client.

## Verifying a client installation using the command line

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the

client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

First, set up the server using the command line, using the instructions in “Setting up the server using the command line.”

Once you have set up the server, you must set up the client, using the instructions in “Connecting an IBM MQ MQI client to a queue manager, using the MQSERVER environment variable” on page 378.

Finally, you can test the communications between client and server, using the instructions in “Testing communication between a client and a server” on page 382.

#### **Related concepts:**

“Verifying a client installation using IBM MQ Explorer” on page 379

You can verify a client installation using the IBM MQ Explorer on Windows and Linux. On the server, you create a queue manager, a local queue, a listener and a server-connection channel. On the client system you create a client-connection channel. Then from the command line you use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

## **Setting up the server using the command line**

Follow these instructions to create a queue manager, queue, and channel on the server. You can then use these objects to verify the installation.

### **About this task**

These instructions assume that no queue manager or other IBM MQ objects have been defined.

IBM MQ object definitions are case-sensitive. Any text entered as an MQSC command in lowercase is converted automatically to uppercase unless you enclose it in single quotation marks. Make sure that you type the examples exactly as shown.

### **Procedure**

1. Create a user ID on the server that is not in the mqm group. This user ID must exist on the server and client. This is the user ID that the sample applications must be run as, otherwise a 2035 error is returned.
2. If your server is running on UNIX, Linux or IBM i, log in as a user in the mqm group.
3. If you are verifying an installation on UNIX, Linux, or Windows, you must set various environment variables so that the installation can be used in the current shell. You can set the environment variables by entering one of the following commands:
  - On UNIX and Linux systems:

```
. MQ_INSTALLATION_PATH/bin/setmqenv -s
```

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.
  - On Windows:

```
. MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.
4. Create a queue manager called QUEUE.MANAGER.1 by entering one of the following commands:

- On UNIX, Linux, and Windows:  
crtmqm QUEUE.MANAGER.1
- On IBM i:  
CRTMQM MQMNAME('QUEUE.MANAGER.1')
- On z/OS, you do not need to create a queue manager.

You see messages telling you that the queue manager has been created.

5. Start the queue manager by entering one of the following commands:

- On UNIX, Linux, and Windows:  
strmqm QUEUE.MANAGER.1
- On IBM i:  
STRMQM MQMNAME('QUEUE.MANAGER.1')
- On z/OS:  
START QMGR

A message tells you when the queue manager has started.

6. On UNIX, Linux, and Windows, start MQSC by entering the following command:

```
runmqsc QUEUE.MANAGER.1
```

A message tells you that an MQSC session has started. MQSC has no command prompt.

7. Define a local queue called QUEUE1 by entering one of the following commands:

- On UNIX, Linux, and Windows:  
DEFINE QLOCAL (QUEUE1)
- On IBM i:  
CRTMQMQ QNAME(QUEUE1) QTYPE(\*LCL)
- On z/OS:  
DEFINE QLOCAL(QUEUE1)

A message tells you when the queue has been created.

8. Allow the user ID that you created in step 1 to use QUEUE1 by entering the following command:

```
SET AUTHREC PROFILE(QUEUE1) OBJTYPE(Queue) PRINCIPAL(' non_mqm_user ') AUTHADD(PUT,GET)
```

where *non\_mqm\_user* is the user ID created in step 1. A message tells you when the authorization has been set. You must also run the following command to give the user ID authority to connect:

```
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL(' non_mqm_user ') AUTHADD(CONNECT)
```

If this command is not run, a 2305 stop error is returned.

9. Define a server-connection channel by entering one of the following commands:

- On all platforms except IBM i:  
DEFINE CHANNEL (CHANNEL1) CHLTYPE (SVRCONN) TRPTYPE (TCP)
- On IBM i:  
CRTMQMCHL CHLNAME(CHANNEL1) CHLTYPE(\*SVRCN) TRPTYPE(\*TCP) MCAUSRID('QMQM')

A message tells you when the channel has been created.

10. Allow your client channel to connect to the queue manager and run under the user ID that you created in step 1, by entering the following MQSC command:

```
SET CHLAUTH(CHANNEL1) TYPE(ADDRESSMAP) ADDRESS(' client_ipaddr ') MCAUSER(' non_mqm_user ')
```

where *client\_ipaddr* is the IP address of the client system, and *non\_mqm\_user* is the user ID created in step 1. A message tells you when the rule has been set.

11. On UNIX, Linux, and Windows, define a listener by entering the following command:



```
DEFINE LISTENER (LISTENER1) TRPTYPE (TCP) CONTROL (QMGR) PORT (port_number)
```

where *port\_number* is the number of the port the listener is to run on. This number must be the same as the number used when defining your client-connection channel in “Installing an IBM MQ client” on page 319 ; see the appropriate HP-UX and Linux sections for more information.

**Note:** If you omit the port parameter from the command, a default value of 1414 is used for the listener port. If you want to specify a port other than 1414, you must include the port parameter in the command, as shown.

12. Start the listener by entering one of the following commands:
  - On UNIX, Linux, and Windows:  
START LISTENER (LISTENER1)
  - On IBM i:  
STRMQMLSR MQMNAME('QUEUE.MANAGER.1') PORT(1414)
  - On z/OS:  
START CHINIT  
START LSTR TRPTYPE(TCP) PORT(*port*)
13. On UNIX, Linux, and Windows, stop MQSC by entering the following command:  
end

You see some messages, followed by the command prompt.

## What to do next

Follow the instructions to set up the client. See “Connecting an IBM MQ MQI client to a queue manager, using the MQSERVER environment variable” on page 378.

### Related concepts:

“Verifying a client installation using the command line” on page 374

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

“Verifying a client installation using IBM MQ Explorer” on page 379

You can verify a client installation using the IBM MQ Explorer on Windows and Linux. On the server, you create a queue manager, a local queue, a listener and a server-connection channel. On the client system you create a client-connection channel. Then from the command line you use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

### Related tasks:

“Testing communication between a client and a server” on page 382

On the IBM MQ MQI client workstation, use the amqsputc sample program to put a message on the queue at the server workstation. Use the amqsgetc sample program to get the message from the queue back to the client.

## Connecting an IBM MQ MQI client to a queue manager, using the MQSERVER environment variable

When an IBM MQ application is run on the IBM MQ MQI client, it requires the name of the MQI channel, the communication type, and the address of the server to be used. Provide these parameters by defining the MQSERVER environment variable.

### Before you begin

Before you start this task, you must complete the task, “Setting up the server using the command line” on page 375, and save the following information:

- The hostname or IP address of the server and port number that you specified when creating the listener.
- The channel name of the server-connection channel.

### About this task

This task describes how to connect an IBM MQ MQI client, by defining the MQSERVER environment variable on the client.

If you are using Windows, HP Integrity NonStop Server, UNIX or Linux systems, you can give the client access to the generated client channel definition table, `amqclchl.tab` instead; see Accessing client-connection channel definitions.

Alternatively, on Windows, if Active Directory support is enabled, the client discovers the client-connection information dynamically from the Active Directory.

### Procedure

1. Log in as the userid that you created in Step 1 of “Setting up the server using the command line” on page 375.
2. Check the TCP/IP connection. From the client, enter one of the following commands:
  - `ping server-hostname`
  - `ping n.n.n.n`

`n.n.n.n` represents the network address. You can set the network address in IPv4 dotted decimal form, for example, `192.0.2.0`. Alternatively, set the address in IPv6 hexadecimal form, for example `2001:0DB8:0204:acff:fe97:2c34:fde0:3485`.

If the **ping** command fails, correct your TCP/IP configuration.

3. Set the MQSERVER environment variable. From the client, enter one of the following commands:
  - a. On Windows:  
`SET MQSERVER=CHANNEL1/TCP/server-address (port)`
  - b. On UNIX and Linux , and IBM MQ client for HP Integrity NonStop Server OSS systems:  
`export MQSERVER=CHANNEL1/TCP/' server-address (port)'`
  - c. On IBM MQ client for HP Integrity NonStop Server Guardian systems:  
`param MQSERVER CHANNEL1/TCP/server-address (port)`
  - d. On IBM i:  
`ADDENVVAR ENVVAR(MQSERVER) VALUE('CHANNEL1/TCP/server-address (port)')`

Where:

- `CHANNEL1` is the server-connection channel name.
- `server-address` is the TCP/IP host name of the server.
- `port` is the TCP/IP port number the server is listening on.

If you do not give a port number, IBM MQ uses the one specified in the `qm.ini` file, or the client configuration file. If no value is specified in these files, IBM MQ uses the port number identified in the TCP/IP services file for the service name `MQSeries`. If an `MQSeries` entry in the services file does not exist, a default value of 1414 is used. It is important that the port number used by the client and the port number used by the server listener program are the same.

## What to do next

Use the sample programs to test communication between the client and server; see “Testing communication between a client and a server” on page 382.

### Related concepts:

“Verifying a client installation using IBM MQ Explorer”

You can verify a client installation using the IBM MQ Explorer on Windows and Linux. On the server, you create a queue manager, a local queue, a listener and a server-connection channel. On the client system you create a client-connection channel. Then from the command line you use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

### Related tasks:

“Setting up the server using the command line” on page 375

Follow these instructions to create a queue manager, queue, and channel on the server. You can then use these objects to verify the installation.

“Testing communication between a client and a server” on page 382

On the IBM MQ MQI client workstation, use the `amqsputc` sample program to put a message on the queue at the server workstation. Use the `amqsgetc` sample program to get the message from the queue back to the client.

## Verifying a client installation using IBM MQ Explorer

You can verify a client installation using the IBM MQ Explorer on Windows and Linux. On the server, you create a queue manager, a local queue, a listener and a server-connection channel. On the client system you create a client-connection channel. Then from the command line you use the sample PUT and GET programs to complete the verification procedure.

To begin the verification setup, see “Setting up the server using IBM MQ Explorer” on page 380.

### Related concepts:

“Verifying a client installation using the command line” on page 374

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

### Related tasks:

“Setting up the client using IBM MQ Explorer” on page 381

You can use MQ Explorer to define the client-connection if you are setting up the client and server on the same workstation on a Windows or Linux system.

“Testing communication between a client and a server” on page 382

On the IBM MQ MQI client workstation, use the amqsputc sample program to put a message on the queue at the server workstation. Use the amqsgetc sample program to get the message from the queue back to the client.

## Setting up the server using IBM MQ Explorer

You can use the IBM MQ Explorer to create a queue manager, queue and server-connection channel on Windows and Linux systems .

### Procedure

1. Create a queue manager:
  - a. Open IBM MQ Explorer.
  - b. Right-click the folder called **Queue Managers**, select **New > Queue Manager**.
  - c. In the first entry field, type the queue manager name, *QUEUE.MANAGER.1*, and click **Finish**.
2. Create a local queue:
  - a. Expand the queue manager you have just created and right-click **queues**.
  - b. Select **New > Local Queue**.
  - c. Enter the queue name, *QUEUE1*, and click **Finish**.
3. Define the server-connection channel:
  - a. Right-click **Channels**.
  - b. Select **New > Server Connection Channel**.
  - c. Enter the channel name, *CHANNEL1*, and click **Next**.
  - d. In the dialog navigation pane, click **MCA** to open the MCA page.
  - e. In the MCA User ID field, enter a userid that is a member of the mqm group, typically your own.
  - f. Click **Finish**.
4. Run the listener.

The listener is automatically started when the queue manager is configured. To check that the listener is running, open **Listeners** and look for LISTENER.TCP.

### What to do next

Set up the client. See “Setting up the client using IBM MQ Explorer” on page 381.

### Related concepts:

“Verifying a client installation using the command line” on page 374

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

### Related tasks:

“Setting up the client using IBM MQ Explorer”

You can use MQ Explorer to define the client-connection if you are setting up the client and server on the same workstation on a Windows or Linux system.

“Testing communication between a client and a server” on page 382

On the IBM MQ MQI client workstation, use the amqsputc sample program to put a message on the queue at the server workstation. Use the amqsgetc sample program to get the message from the queue back to the client.

## Setting up the client using IBM MQ Explorer

You can use MQ Explorer to define the client-connection if you are setting up the client and server on the same workstation on a Windows or Linux system.

### Procedure

1. Select the queue manager, *QUEUE.MANAGER.1*
2. Open the **Channels** folder, then right-click **Client Connections** > **New** > **Client-connection Channel...**
3. Enter the channel name, *CHANNEL1*, for the client connection, and click **Next**.
4. Enter the queue manager name, *QUEUE.MANAGER.1*
5. Enter the following string as the connection name:  
*server-address (port)*

Where:

- *server-address* is the TCP/IP host name of the server
  - *port* is the TCP/IP port number the server is listening on
6. Click **Finish**.
  7. From the command line, set the MQCHLLIB environment variable:
    - For Windows clients, enter the following command, where *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed:  

```
SET MQCHLLIB= MQ_INSTALLATION_PATH\qmgrs\QUEUE!MANAGER!1\@ipcc
```
    - For Linux clients, enter the following command:  

```
export MQCHLLIB=var/mqm/qmgrs/QUEUE!MANAGER!1/@ipcc
```

**Note:** The queue manager name contains “.”. IBM MQ creates the queue manager directory with the name, *QUEUE!MANAGER!1*

### What to do next

Use the sample programs to test communication between the client and server. See “Testing communication between a client and a server” on page 382.

### Related concepts:

“Verifying a client installation using the command line” on page 374

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

### Related tasks:

“Testing communication between a client and a server”

On the IBM MQ MQI client workstation, use the amqsputc sample program to put a message on the queue at the server workstation. Use the amqsgetc sample program to get the message from the queue back to the client.

“Setting up the server using IBM MQ Explorer” on page 380

You can use the IBM MQ Explorer to create a queue manager, queue and server-connection channel on Windows and Linux systems .

## Testing communication between a client and a server

On the IBM MQ MQI client workstation, use the amqsputc sample program to put a message on the queue at the server workstation. Use the amqsgetc sample program to get the message from the queue back to the client.

### Before you begin

Complete the previous topics in this section:

- Set up a queue manager, channels, and queue.
- Open a command window.
- Set system environment variables.

### About this task

Note that IBM MQ object definitions are case-sensitive. Text entered as an MQSC command in lowercase is converted automatically to uppercase unless you enclose it in single quotation marks. Make sure that you type the examples exactly as shown.

### Procedure

1. On UNIX and Linux systems, change into the *MQ\_INSTALLATION\_PATH*/samp/bin directory, which contains the sample programs. For IBM MQ client for HP Integrity NonStop Server, change into the *MQ\_INSTALLATION\_PATH*/opt/mqm/samp/bin directory, which contains the sample programs. On Windows systems, change into the *MQ\_INSTALLATION\_PATH*\Tools\C\Samples\Bin directory for 32 bit systems or the *MQ\_INSTALLATION\_PATH*\Tools\C\Samples\Bin64 directory for 64 bit systems. *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.
2. If you are verifying an installation on a UNIX, Linux, or Windows system, you must set certain environment variables so that the installation can be used in the current shell. This step is not applicable to IBM MQ client for HP Integrity NonStop Server. You can set the environment variables by entering one of the following commands:

- Windows:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.

- UNIX and Linux:  
`. MQ_INSTALLATION_PATH/bin/setmqenv -s`

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.

3. Start the PUT program for QUEUE1 on QUEUE.MANAGER.1 by entering one of the following commands:

- Windows:  
`amqsputc QUEUE1 QUEUE.MANAGER.1`
- UNIX and Linux, and IBM MQ client for HP Integrity NonStop Server :  
`./amqsputc QUEUE1 QUEUE.MANAGER.1`

If the command is successful, the following messages are displayed:

```
Sample AMQSPUT0 start target queue is QUEUE1
```

**Tip:** You might get the error, MQRC\_NOT\_AUTHORIZED ( 2035 ). By default, channel authentication is enabled when a queue manager is created. Channel authentication prevents privileged users accessing a queue manager as an IBM MQ MQI client. For verifying the installation, you can either change the MCA user ID to a non-privileged user, or disable channel authentication. To disable channel authentication run the following MQSC command:

```
ALTER QMGR CHLAUTH(DISABLED)
```

When you finish the test, if you do not delete the queue manager, re-enable channel authentication:

```
ALTER QMGR CHLAUTH(ENABLED)
```

4. Type some message text, then press **Enter** twice. The following message is displayed:

```
Sample AMQSPUT0 end
```

Your message is now on the queue that is on the server queue manager.

5. Start the GET program for QUEUE1 on QUEUE.MANAGER.1 by entering one of the following commands:

- On Windows:  
`amqsgetc QUEUE1 QUEUE.MANAGER.1`
- On UNIX and Linux, and IBM MQ client for HP Integrity NonStop Server:  
`./amqsgetc QUEUE1 QUEUE.MANAGER.1`
- On IBM i:  
`CALL PGM(QMQM/AMQSGETC) PARM(QUEUE1 QUEUE.MANAGER.1)`

The sample program starts, and your message is displayed. After a short pause (approximately 30 seconds), the sample ends and the command prompt is displayed again.

## Results

You have now successfully verified the client installation.

## What to do next

1. If your server is on UNIX, Linux, or Windows, you must set various environment variables so that the installation can be used in the current shell. You can set the environment variables by entering one of the following commands:

- On UNIX and Linux:  
`. MQ_INSTALLATION_PATH/bin/setmqenv -s`

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.

- On Windows:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.

2. On the server, stop the queue manager by entering one of the following commands:

- On UNIX, Linux, and Windows:

```
endmqm QUEUE.MANAGER.1
```

- On IBM i:

```
ENDMQM MQMNAME(QUEUE.MANAGER.1)
```

- On z/OS:

```
STOP CHINIT
STOP QMGR
```

3. On the server, delete the queue manager by entering one of the following commands:

- On UNIX, Linux, and Windows:

```
dltmqm QUEUE.MANAGER.1
```

- On IBM i:

```
DLTMQM MQMNAME(QUEUE.MANAGER.1)
```

- On z/OS, you do not need to delete the queue manager.

#### Related concepts:

“Verifying a client installation using the command line” on page 374

You can verify a client installation using the command line. On the server you create a queue manager, a local queue, a listener, and a server-connection channel. You must also apply security rules to allow the client to connect and make use of the queue defined. On the client you create a client-connection channel, and then use the sample PUT and GET programs to complete the verification procedure.

“Verifying a client installation using IBM MQ Explorer” on page 379

You can verify a client installation using the IBM MQ Explorer on Windows and Linux. On the server, you create a queue manager, a local queue, a listener and a server-connection channel. On the client system you create a client-connection channel. Then from the command line you use the sample PUT and GET programs to complete the verification procedure.

“Installing an IBM MQ client” on page 319

“Verifying a server installation” on page 363

You can verify a local (stand-alone) installation or a server-to-server installation of the IBM MQ server. A local installation has no communication links with other IBM MQ installations while a server-to-server installation does have links to other installations.

#### Related tasks:

“Installing an IBM MQ server” on page 254

After preparing your system for installation you can install IBM MQ by following the appropriate instructions for your platform. After installation, you might want to verify your installation to check that installation has been successful.

---

## Verifying the installation of IBM MQ Telemetry

There are three ways to verify the installation of IBM MQ Telemetry. Any can be used, regardless of whether IBM MQ Telemetry was installed as a custom installation of IBM MQ, or added to an existing installation of IBM MQ.

Within IBM MQ you can verify the installation of IBM MQ Telemetry in either of the following ways:

- “Verifying the installation of IBM MQ Telemetry by using MQ Explorer” on page 385
- “Verifying the installation of IBM MQ Telemetry using the command line” on page 387



You can also verify the installation by using the MQTT messaging client for JavaScript in a browser that supports the RFC 6455 (WebSocket) standard. A version of this client is installed with IBM MQ Telemetry, and the latest version is available as part of the free download IBM Messaging Telemetry Clients SupportPac. The steps for using this client with IBM MQ Telemetry are given in the client pack information set, in *Getting started with the MQTT messaging client for JavaScript*. To verify the IBM MQ Telemetry installation you do not need the latest version of the client, so you should omit step 1.

## Verifying the installation of IBM MQ Telemetry by using MQ Explorer

Use the Define sample configuration wizard and the MQTT client utility in MQ Explorer to verify that the IBM MQ Telemetry components have installed. Also check that publish/subscribe works correctly.

### Before you begin

The IBM MQ Telemetry runtime and support for MQ Explorer must be installed. The telemetry folder is part of a queue manager. To view the telemetry folder, you must start a queue manager.

Before running the define sample configuration wizard on an existing queue manager, review the information provided by the wizard about the configuration changes that are made. The changes might have implications for the configuration of the existing queue manager. Alternatively, run the sample configuration wizard on a newly created queue manager to avoid changing any security settings.

### About this task

To configure IBM MQ Telemetry there is a define sample configuration wizard that can be run from MQ Explorer. The wizard runs through a series of steps, including defining and starting the telemetry (MQXR) service, setting up the default transmission queue, and configuring a telemetry channel.

If you would prefer to do this manually, see *Configuring a queue manager for telemetry on Linux and AIX*. For Windows, see *Configuring a queue manager for telemetry on Windows*.

You can open the define sample configuration wizard from the IBM MQ Telemetry Welcome page in MQ Explorer. The wizard determines which steps are required based on the current configuration.

For example, the following actions might be specified by the wizard:

- Define the telemetry (MQXR) service.
- Start the telemetry (MQXR) service.
- Define the telemetry transmit queue.
- Set the default transmit queue of the queue manager to `SYSTEM.MQTT.TRANSMIT.QUEUE`.

If telemetry is already configured for this queue manager, the link to open the wizard is replaced with static text. The text confirms that the sample configuration has been set up.

After the configuration has finished, you can use MQ Explorer to open the MQTT client utility. Use the MQTT client utility to verify that IBM MQ Telemetry is set up correctly.

The following items summarize the main goals that can be achieved using the MQTT client utility:

- Validation of a basic or custom IBM MQ Telemetry configuration by connecting, subscribing to topics and publishing messages.
- Showcases the main features of MQTT protocol.
- Provides a simple tool to aid in debugging IBM MQ Telemetry applications.

You can find additional information within the MQ Explorer by using the **Help** menu or pressing the **F1** key.

## Procedure

1. Start MQ Explorer.

On Windows and Linux systems, you can start MQ Explorer by using the system menu, the MQExplorer executable file, the **mqexplorer** command, or the **strmqcfcg** command.

2. Open the Welcome to MQ Telemetry page.

- To use an existing queue manager, click on IBM MQ\Queue Managers\*qMgrName*\Telemetry folder to open the Welcome to MQ Telemetry page.
- If, for the reasons mentioned, you decide to use a new queue manager,
  - a. Click **Queue Managers > New > Queue Manager**.
  - b. Type MQTTVerification as the **Queue manager name > Next > Next > Next**.
  - c. Change the default port in **Listen on port number**, if the port is in use > **Finish**.
  - d. When the queue manager starts, click on IBM MQ\Queue Managers\MQTTVerification\Telemetry folder to open the Welcome to MQ Telemetry page.

3. From the Welcome to MQ Telemetry page in MQ Explorer, click **Define sample configuration**.

If this link is not present, and instead you see the text, “The sample configuration has been set up for this queue manager”, then telemetry has already been configured. Proceed to step 6.

If you clicked **Define sample configuration**, the page opens, and lists actions that are to be performed as part of the sample configuration.

4. Leave **Launch MQTT client utility** checked, if you want to automatically start the MQTT client utility. The check box is selected by default.
5. Click **Finish**.
6. Click **Connect**.

In the MQTT client utility panel, ensure that the host and port names are correct.

If you did not automatically start the MQTT client utility panel in step 4, you can start it either by using a direct link from the Welcome to MQ Telemetry panel, or by right-clicking a NON-SSL channel, which allows you to control the channel it runs on.

The client history records a Connected event.

7. Click **Subscribe**.

The client history records a Subscribed event.

8. Click **Publish**.

The client history records a Published and Received event.

## Results

If the publish/subscribe finishes successfully, the IBM MQ Telemetry installation is verified.

If you encounter problems during the installation process, view the error log:

- On Windows, the default location for this log is, *WebSphere MQ data directory*\qmgrs\*qMgrName*\mqxr
- On AIX and Linux, the default location for this log is, */var/mqm/qmgrs/qMgrName/mqxr/*

# Verifying the installation of IBM MQ Telemetry using the command line

Follow these instructions to run scripts and a sample application to verify that the IBM MQ Telemetry components have installed, and are able to publish and subscribe.

## Before you begin

The telemetry (MQXR) service must be started to run the sample programs. The user ID must be a member of the `mqm` group.

The `SampleMQM` script creates and uses a queue manager called `MQXR_SAMPLE_QM`. Therefore, do not run unaltered on a system that already has a `MQXR_SAMPLE_QM` queue manager. Any changes made might have implications for the configuration of the existing queue manager.

This task uses the `MQTTV3` sample Java application, and the associated Java client library. These resources are no longer included in IBM MQ Telemetry, they are now part of the free download IBM Messaging Telemetry Clients SupportPac. You install this pack as part of this task.

There are two commands to run the `MQTTV3` sample Java application. The first command creates a subscription, then waits for a message. The second command publishes to that subscription. Therefore the commands must be entered into different command lines or shell windows.

## About this task

To perform verification on a server or device without a GUI, scripts are provided in the `samples` directory. The `SampleMQM` script performs the required steps to configure IBM MQ Telemetry. The `MQTTV3` sample Java application can then be run to validate the basic or custom IBM MQ Telemetry configuration by connecting, subscribing to topics, and publishing messages. The `CleanupMQM` sample script can be run to delete the queue manager created by the `SampleMQM` script.

The following items summarize the main goals that can be achieved using this verification procedure:

- Validate a basic or custom IBM MQ Telemetry configuration by connecting, subscribing to topics and publishing messages.
- Showcase the main features of the MQTT protocol.
- Provide a simple tool to aid in debugging IBM MQ Telemetry applications.

## Procedure

1. Download the free IBM Messaging Telemetry Clients SupportPac.
2. Uncompress the client pack into a directory of your own choosing.

The sample applications and client libraries are in client-specific directories under `<CLIENTPACKDIR>/SDK/clients`, where `<CLIENTPACKDIR>` is the directory in which you uncompressed the client pack. This task uses the `MQTTV3` sample Java application, and the associated Java client library. These resources are available in the following location:

```
<CLIENTPACKDIR>/SDK/clients/java
```

3. Configure IBM MQ Telemetry.

The `SampleMQM` script runs through a series of steps, including creating the `MQXR_SAMPLE_QM` queue manager, defining and starting the telemetry (MQXR) service, setting up the default transmission queue, and configuring a telemetry channel.

For information about performing this manually, see [Configuring a queue manager for telemetry on Linux and AIX](#), or [Configuring a queue manager for telemetry on Windows](#).

- On Windows systems, enter the following command in a command line:  

```
<MQINSTDIR>\mqxr\samples\SampleMQM.bat
```
- On AIX or Linux systems, enter the following command in a shell window:

```
<MQINSTDIR>/mqxr/samples/SampleMQM.sh
```

where <MQINSTDIR> is the install directory for this installation of IBM MQ. A queue manager called MQXR\_SAMPLE\_QM is created, and IBM MQ Telemetry is configured.

4. Run the MQTTV3 sample Java application to create a subscription.
  - On Windows systems, enter the following commands in a command line:

```
java -cp
"<CLIENTPACKDIR>\SDK\clients\java\org.eclipse.paho.sample.mqttv3app.jar;
<CLIENTPACKDIR>\SDK\clients\java\org.eclipse.paho.client.mqttv3.jar"
org.eclipse.paho.sample.mqttv3app.Sample -a subscribe
```

- On AIX or Linux systems, enter the following commands in a shell window:

```
java -cp
<CLIENTPACKDIR>/SDK/clients/java/org.eclipse.paho.sample.mqttv3app.jar:
<CLIENTPACKDIR>/SDK/clients/java/org.eclipse.paho.client.mqttv3.jar
org.eclipse.paho.sample.mqttv3app.Sample -a subscribe
```

The subscription is created, and waits to receive a message.

5. Run the MQTTV3 sample Java application to publish to the subscription.
  - On Windows systems, enter the following command in a second command line:

```
java -cp
"<CLIENTPACKDIR>\SDK\clients\java\org.eclipse.paho.sample.mqttv3app.jar;
<CLIENTPACKDIR>\SDK\clients\java\org.eclipse.paho.client.mqttv3.jar"
org.eclipse.paho.sample.mqttv3app.Sample -m "Hello from an MQTT v3 application"
```

- On AIX or Linux systems, enter the following command in a second shell window:

```
java -cp
<CLIENTPACKDIR>/SDK/clients/java/org.eclipse.paho.sample.mqttv3app.jar:
<CLIENTPACKDIR>/SDK/clients/java/org.eclipse.paho.client.mqttv3.jar
org.eclipse.paho.sample.mqttv3app.Sample -m "Hello from an MQTT v3 application"
```

The message Hello from an MQTT v3 application, that you typed into the second command line or shell window, is published by that application and received by the application in the first window. The application in the first window shows it on the screen.

6. Press **Enter** in the first command line or shell window to end the subscribing application.
7. Remove the queue manager created by the SampleMQM script.
  - On Windows systems, enter the following command in a command line:  
<MQINSTDIR>\mqxr\samples\CleanupMQM.bat
  - On AIX or Linux systems, enter the following command in a shell window:  
<MQINSTDIR>/mqxr/samples/CleanupMQM.sh

## Results

If the scripts finished, and messages can be sent and received, the IBM MQ Telemetry installation is verified.

## What to do next

If you encounter any problems during the verification process, see IBM MQ Telemetry troubleshooting. You can also view the error log:

- On Windows systems, the default location for the queue manager log is <MQINSTDIR>\qmgrs\MQXR\_SAMPLE\_QM\mqxr
- On AIX and Linux systems, the default location for the queue manager log is /var/mqm/qmgrs/MQXR\_SAMPLE\_QM/mqxr/

---

## Uninstalling IBM MQ components

Information to help you uninstall components of IBM MQ.

### About this task

Refer to the following topics to uninstall components of IBM MQ:


- Uninstalling IBM MQ server and client
- Uninstalling IBM MQ Telemetry
- Uninstalling IBM MQ Advanced Message Security

---

## Uninstalling

The topics in this section provide instructions on how to uninstall components.

Select the appropriate topic for your platform to find out how to uninstall IBM MQ components:

- “Uninstalling IBM MQ on AIX”
- “Uninstalling IBM MQ on HP-UX” on page 391
- “Uninstalling IBM MQ on Linux” on page 392
- “Uninstalling IBM MQ on Solaris” on page 394
-  “Uninstalling IBM MQ for IBM i” on page 400
- “Uninstalling IBM MQ on Windows systems” on page 395

### Related concepts:

“Installing IBM MQ” on page 249

The topics in this section provide instructions on how to install IBM MQ.

“Verifying an IBM MQ installation” on page 363

The topics in this section provide instructions on how to verify a server or a client installation of IBM MQ on distributed systems.

## Uninstalling IBM MQ on AIX

On AIX, you can uninstall the IBM MQ server or client using the System Management Interface Tool (SMIT) or the `installp` command.

### Before you begin

If any updates have been applied, remove them before starting this uninstallation procedure. For more information, see AIX: Restoring the previous maintenance level on IBM WebSphere MQ Version 7.5 .

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group `mqm`.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.

- c. Display the state of all queue managers on the system. Enter the following command:  
`dspmqr -o installation`
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:  
`endmqm QMgrName`
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:  
`endmqm -m QMgrName`
3. Log in as root.
  4. Uninstall IBM MQ using either **installp** or **smitt**. If IBM MQ was installed in a non-default location, you must use **installp** to uninstall.
    - Uninstall using **installp** by entering one of the following commands:
      - For an installation in the default location `/usr/mqm`  
`installp -u mqm`
      - For an installation in a non-default location:  
`installp -R usil -u mqm`

where *usil* is the path of the User Specified Install Location (USIL) specified when the product was installed.
    - Uninstall using **smitt**:
      - a. Select the required **smitt** window using the following sequence:  
 Software Installation and Maintenance  
 Software Maintenance and Utilities  
 Remove Installed Software
      - b. List the software in the **SOFTWARE name** field:
        - 1) Enter `.`
        - 2) Press **F4**
      - c. Select the file sets to uninstall from the list (those beginning with `mqm`), and press **Enter**. There is an option at this stage to do a preview. Leave the option set to the default value of **Yes** to preview the file sets you are uninstalling, or select **No** to not preview these file sets.
      - d. Press **Enter** on the **Remove Installed Software** panel, it asks whether you are sure, press **Enter**.

## Results

After uninstallation, certain files under the directory trees `/var/mqm` and `/etc/opt/mqm` are not removed. These files contain user data and remain so subsequent installations can reuse the data. Most of the remaining files contain text, such as INI files, error logs, and FDC files. The directory tree `/var/mqm/shared` contains files that are shared across installations, including the executable shared libraries `libmqzsd.a` and `libmqzsd_r.a`.

## What to do next

- If the product successfully uninstalled, you can delete any files and directories contained in the `/usr/mqm` directory under the User Specified Install Location (USIL) specified in the **installp** uninstallation command.
- Use the **lspp** command to check for other products installed in the USIL. If there are no other products installed in the USIL and you do not intend to use it again, you can delete the USIL using the **rmusil** command.
- If there are no other IBM MQ installations on the system, and you are not planning to reinstall or migrate, you can delete the `/var/mqm` and `/etc/opt/mqm` directory trees, including the files `libmqzsd.a` and `libmqzsd_r.a`. Deleting these directories destroys all queue managers and their associated data.

## Uninstalling IBM MQ on HP Integrity NonStop Server

On HP Integrity NonStop Server systems, you can uninstall the IBM MQ client by using the **rm** command.

### Procedure

1. Stop all IBM MQ applications that are associated with the installation you are uninstalling.
2. Log in to the OSS as the user ID that owns the installation.
3. Use the OSS **rm** command to delete the files from the Guardian subvolume used by the installation. For example, use the following command:

```
rm -rf <mqpath>/opt/mqm/bin/G/*
```

4. Use the OSS **rm** command to delete the OSS directory trees for the installation. For example, use the following command:

```
rm -rf <mqpath>
```

## Uninstalling IBM MQ on HP-UX

On HP-UX, you can uninstall the IBM MQ server or client using the **swremove** command.

### Before you begin

If any updates have been applied, remove them before starting this uninstallation procedure. For more information, see HP-UX: Restoring the previous maintenance level on IBM WebSphere MQ Version 7.5 .

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group `mqm`.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:

```
dspmqr -o installation
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:

```
endmqm QMgrName
```
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:

```
endmqlsr -m QMgrName
```
3. Log in as root.
4. Uninstall IBM MQ using **swremove**:
  - To uninstall all IBM MQ components, enter the following command:

```
swremove MQSERIES,1= MQ_INSTALLATION_PATH
```

where `MQ_INSTALLATION_PATH` is the path where IBM MQ is installed.
  - To uninstall selected IBM MQ components, enter the following command:

```
swremove componentname,1= MQ_INSTALLATION_PATH
```

where *componentname* is the name of the component to uninstall, and *MQ\_INSTALLATION\_PATH* is the path where IBM MQ is installed.

For example:

- To uninstall the client component, from an installation in */opt/myLocation*, enter the following command:

```
swremove MQSERIES.MQM-CL-HPUX,1=/opt/myLocation
```

- To uninstall the client and the telemetry client components, from an installation in */opt/myLocation*, enter the following command:

```
swremove MQSERIES.MQM-CL-HPUX,1=/opt/myLocation MQSERIES.MQM-TXCLIENT,1=/opt/myLocation
```

## Results

After uninstallation, certain files under the directory trees */var/mqm* and */etc/opt/mqm* are not removed. These files contain user data and remain so subsequent installations can reuse the data. Most of the remaining files contain text, such as INI files, error logs, and FDC files. The directory tree */var/mqm/shared* contains files that are shared across installations, including the executable shared libraries *libmqzsd.so* and *libmqzsd\_r.so*.

## What to do next

- If the product successfully uninstalled, you can delete any files and directories contained in the installation directory.
- If there are no other IBM MQ installations on the system, and you are not planning to reinstall or migrate, you can delete the */var/mqm* and */etc/opt/mqm* directory trees, including the files *libmqzsd.so* and *libmqzsd\_r.so*. Deleting these directories destroys all queue managers and their associated data.

## Uninstalling IBM MQ on Linux

On Linux, you can uninstall the IBM MQ server or client using the **rpm** command.

### Before you begin

If you have applied one or more fix packs to the version of IBM MQ that you want to uninstall, you need to remove the fix packs in reverse chronological install order before you remove the base packages.

You must remove any updates before starting this uninstallation procedure. For more information, see [Linux: Restoring the previous maintenance level on IBM MQ](#).

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group *mqm*.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv -s
```

where *MQ\_INSTALLATION\_PATH* refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:

```
dspmqr -o installation
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:

```
endmqm QMGrName
```



- e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:

```
endmq1sr -m QMgrName
```

3. Log in as root.

4. Uninstall IBM MQ using the **rpm** command:

- a. On a system with a single installation:

- 1) Find out the names of the packages (components) currently installed on your system, by entering the following command:

```
rpm -qa | grep MQSeries
```

- 2) Remove all the components at the same time by appending all the package names to the **rpm** command arguments. For example, to remove the runtime, Server and SDK components enter the following command:

```
rpm -ev MQSeriesRuntime MQSeriesServer MQSeriesSDK
```

**Tip:** To list the packages, and uninstall them in one go, use a command like the following one:

```
rpm -qa | grep MQSeries | xargs rpm -ev
```

- 3) If you are using Ubuntu, add the **--force-debian** attribute. For example, to remove the runtime, Server and SDK components enter the following command:

```
rpm --force-debian -ev MQSeriesRuntime MQSeriesServer MQSeriesSDK
```

- b. On a system with multiple installations:

- 1) Find out the names of the packages (components) currently installed on your system, by entering the following command:

```
rpm -qa | grep suffix
```

where *suffix* is the unique name given to the packages when **crtmqpkg** was run at installation time. *suffix* is included in each of the package names that belong to a particular installation.

- 2) Remove all the components at the same time by appending all the package names to the **rpm** command arguments. For example, to remove the runtime, Server and SDK components for an installation with the *suffix* MQ80, enter the following command:

```
rpm -ev MQSeriesRuntime-MQ80 MQSeriesServer-MQ80 MQSeriesSDK-MQ80
```

- 3) If you are using Ubuntu, add the **--force-debian** attribute. For example, to remove the runtime, Server and SDK components for an installation with the *suffix* MQ80, enter the following command:

```
rpm --force-debian -ev MQSeriesRuntime-MQ80 MQSeriesServer-MQ80 MQSeriesSDK-MQ80
```

## Results

After uninstallation, certain files under the directory trees `/var/mqm` and `/etc/opt/mqm` are not removed. These files contain user data and remain so subsequent installations can reuse the data. Most of the remaining files contain text, such as INI files, error logs, and FDC files. The directory tree `/var/mqm/shared` contains files that are shared across installations, including the executable shared libraries `libmqzsd.so` and `libmqzsd_r.so`.

## What to do next

- If the product successfully uninstalled, you can delete any files and directories contained in the installation directory.
- If there are no other IBM MQ installations on the system, and you are not planning to reinstall or migrate, you can delete the `/var/mqm` and `/etc/opt/mqm` directory trees, including the files `libmqzsd.so` and `libmqzsd_r.so`. Deleting these directories destroys all queue managers and their associated data.

## Uninstalling IBM MQ on Solaris

On Solaris, you can uninstall the IBM MQ server or client using the **pkgrm** command.

### Before you begin

If any updates have been applied, remove them before starting this uninstallation procedure. For more information, see Solaris: Restoring the previous maintenance level on IBM MQ .

**Restriction:** On Solaris, you cannot remove components from an install. There is no supported method of doing this.

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group `mqm`.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:

```
dspmqr
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:

```
endmqm QMgrName
```
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:

```
endmqrsr -m QMgrName
```
3. Log in as root.
4. Uninstall IBM MQ using **pkgrm**:
  - a. On a system with a single installation, enter the following command:

```
pkgrm mqm
```
  - b. On a system with multiple installations:

```
pkgrm mqm-suffix
```

where *suffix* is the unique name given to the packages when **crtmqpkg** was run at installation time. *suffix* is included in each of the package names that belong to a particular installation. The first installation on the system does not have a *suffix*, and is uninstalled using the same method as for a single installation.

If a package has a dependency on `mqm`, **pkgrm** returns the name of the package. Uninstall the dependent packages first.

### Results

After uninstallation, certain files under the directory trees `/var/mqm` and `/etc/opt/mqm` are not removed. These files contain user data and remain so subsequent installations can reuse the data. Most of the remaining files contain text, such as INI files, error logs, and FDC files. The directory tree `/var/mqm/shared` contains files that are shared across installations, including the executable shared library `libmqzsd.so`.

## What to do next

- If the product successfully uninstalled, you can delete any files and directories contained in the installation directory.
- If there are no other IBM MQ installations on the system, and you are not planning to reinstall or migrate, you can delete the `/var/mqm` and `/etc/opt/mqm` directory trees, including the file `libmqzsd.so`. Deleting these directories destroys all queue managers and their associated data.

## Uninstalling IBM MQ on Windows systems

You can uninstall the IBM MQ MQI clients and servers on Windows systems by using the control panel, the command line ( `msiexec` ), `MQParms`, or by using the installation media, in which case you can optionally remove queue managers as well.

### Before you begin

By default, uninstall logging is not enabled on Windows. To ensure that you receive an uninstall log, carry out the following procedure:

1. In a command prompt, open the registry editor by issuing the command **regedit**.
2. Go to the appropriate registry key: `HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Installer`
3. Under this registry key add the following information:

**Name** Logging

**Data type**  
REG\_SZ

**Value** voicewarmup

4. Save the updated registry key.

### Procedure

The first part of the procedure ensures that there are no IBM MQ programs or processes running:

1. If you are running IBM MQ with the Microsoft Cluster Service (MSCS), remove the queue managers from MSCS control before uninstalling IBM MQ. Perform the following steps for each queue manager currently under MSCS control :
  - a. Take the queue manager resource offline.
  - b. Destroy the resource instance.
  - c. Migrate the queue manager files back from shared drives. This step is shown as optional in Removing a queue manager from MSCS control. However, it is mandatory in this case.
2. Stop all IBM MQ applications associated with the installation you are uninstalling.
3. Close all IBM MQ Managed File Transfer agents. If you have an IBM MQ Managed File Transfer agent running, close it by using the **fteStopAgent** command; see `fteStopAgent` (stop an IBM MQ Managed File Transfer agent).
4. For a server installation, end all IBM MQ activity:
  - a. Log in as a user in the group `mqm`.
  - b. Stop all running queue managers and listeners by using the IBM MQ Explorer, or by entering the following commands:
    - 1) Set up your environment to work with the installation you want to uninstall by entering the following command:

```
MQ_INSTALLATION_PATH\bin\setmqenv -s
```

where `MQ_INSTALLATION_PATH` is the location where IBM MQ is installed.
    - 2) For each queue manager, enter the following command to stop the queue manager:

```
endmqm queue_manager_name
```

- 3) For each queue manager, enter the following command to stop any listeners associated with the queue manager:

```
endmq1sr -m queue_manager_name
```

5. Stop IBM MQ. To do this right-click the **IBM MQ** icon in the system tray, then select **Stop IBM MQ**.
6. Close all IBM MQ windows.
7. Stop any monitoring service.

When all processes associated with IBM MQ are no longer running, you can uninstall IBM MQ:

8. Uninstall IBM MQ by using one of the following methods:
  - Use the Windows Control Panel. This process is described in: “Uninstalling IBM MQ by using the control panel” on page 397. This method does not remove the queue manager data.
  - Use the command line by running the **msiexec** command as described in: “Uninstalling IBM MQ by using the command line” on page 397. This method does not remove the queue manager data.
  - Use the appropriate parameters with **MQParms**. This process is described in “Uninstalling IBM MQ by using MQParms” on page 399. This method does not remove the queue manager data.
  - Use the installation media, by selecting the appropriate option as described in: “Uninstalling IBM MQ on Windows by using the installation media” on page 399. The option to remove queue manager data is displayed in the Removing Server feature panel, if appropriate.

If you have to cancel the uninstall process before it is finished, you might have to reconfigure IBM MQ with the Prepare IBM MQ wizard because the rollback of the deletion of the IBM MQ service is unable to set the service's user account password. Use the following command to reconfigure IBM MQ:

```
MQ_INSTALLATION_PATH\bin\amqmjpse.exe -r
```

For more information about the Prepare IBM MQ wizard, see “Configuring IBM MQ with the Prepare IBM MQ wizard” on page 295.

9. Check the Windows event log and restart the system if necessary. If event ID 10005 is written to the Windows event log, you must restart the system to complete the uninstall.
10. Optional: If you are uninstalling the last or only installation of IBM MQ, you can remove all the information about previous installations that is retained on the system, if you want to.

Two registry values remain after uninstallation:

- 32 bit systems:
  - My Computer\HKEY\_LOCAL\_MACHINE\SOFTWARE\IBM\WebSphere MQ\LogDefaultPath
  - My Computer\HKEY\_LOCAL\_MACHINE\SOFTWARE\IBM\WebSphere MQ\WorkPath
- 64 bit systems:
  - My Computer\HKEY\_LOCAL\_MACHINE\SOFTWARE\WOW6432Node\IBM\WebSphere MQ\LogDefaultPath
  - My Computer\HKEY\_LOCAL\_MACHINE\SOFTWARE\WOW6432Node\IBM\WebSphere MQ\WorkPath

Data folders will also remain and are located at *MQ\_DATA\_PATH*\Config, where *MQ\_DATA\_PATH* is the location of the IBM MQ data directory. Most of the remaining files contain text such as INI files, error logs, and FDC files. The executable shared library mqzsd.dll also remains.

If a client is installed on a system where the LogDefaultPath registry value remains from a previous server install, a client installation will attempt to create this directory if it does not already exist. If this behavior is not wanted, remove the LogDefaultPath registry value before installing the client.

## Uninstalling IBM MQ by using the control panel

You can uninstall IBM MQ by using the control panel to remove all currently installed features.

### Before you begin

Start the uninstalling process by following the steps described in “Uninstalling IBM MQ on Windows systems” on page 395.

If you no longer require the queue managers that are on the system, delete them by using the IBM MQ Explorer or the `dlmqm` command.

### Procedure

1. From the Windows taskbar, open the control panel by clicking **Start > Settings > Control Panel**, or **Start > Control Panel**.
2. Open **Programs and Features**.
3. Click **IBM MQ (*installation\_name*)**, where *installation\_name* is the name of the installation you want to remove.
4. Click **Remove** or **Uninstall** and click **Yes** to confirm. If User Account Control (UAC) is enabled, accept the Windows prompt to allow the uninstallation to run as elevated. The program then begins and runs to completion.

### What to do next

Complete the steps that you started in “Uninstalling IBM MQ on Windows systems” on page 395.

## Uninstalling IBM MQ by using the command line

You can uninstall IBM MQ by running the `msiexec` command from the command line to remove all currently installed features.

### Before you begin

Start the uninstalling process by following the steps described in “Uninstalling IBM MQ on Windows systems” on page 395.

If you no longer require the queue managers that are on the system, delete them by using the IBM MQ Explorer or the `dlmqm` command.

### About this task

To start uninstalling, use the `msiexec` command.

If you are running IBM MQ on Windows with User Account Control (UAC) enabled, you must invoke the silent uninstallation from an elevated command prompt. Elevate a command prompt by using a right-click to start the command prompt and choose **Run as administrator**.

In all of the examples of commands shown, the variable names used are as follows:

- *installation\_name* is the name of the installation you want to remove.
- *product\_code* is the value shown for MSIProdCode in the output of the following command:  
`dspmqinst -n installation_name`

An example of a product code is {0730749B-080D-4A2E-B63D-85CF09AE0EF0}.

- *response\_file* is the file that contains the [Response] stanza and the required *property = value* pairs. For details about how to create a response file, see “Creating a response file” on page 287. For details

of the parameters you can specify in a response file, see Table 48 on page 282 in “Advanced installation using `msiexec`” on page 279. This is an example of a simple uninstallation [Response] stanza:

```
[Response] REMOVE="ALL"
```

## Procedure

To uninstall all IBM MQ features, use one of the following methods:

- Run the `msiexec` command with a parameter that calls a response file.

A response file is an ASCII text file that contains the parameter values that you want to set for the uninstallation. The response file has a format similar to a Windows `.ini` file, and contains the stanza [Response]. This stanza contains parameters that the `msiexec` command can use, in the form of *property = value* pairs. The `msiexec` command ignores any other stanzas in the file.

You can set which features to uninstall, and set whether to keep existing queue managers.

To silently uninstall IBM MQ using a response file, enter the following command:

```
msiexec /i {product_code} /! *v "c:\removal.log" /q USEINI="response_file" INSTALLATIONNAME="installation_name"
```

- Enter one of the following commands on the command line:

- To invoke an interactive uninstallation giving you the option to remove queue manager data (providing there are no other IBM MQ installations remaining):

```
msiexec /i {product_code} /! *v "c:\removal.log" REMOVE="All" INSTALLATIONNAME="installation_name"
```

If you are running IBM MQ on a Windows system with User Account Control (UAC) enabled, you might see Open File - Security Warning dialog boxes during uninstallation that list International Business Machines Limited as the publisher. Click **Run** to allow the uninstallation to continue.

- To invoke a silent uninstallation that does not remove any queue manager data:

```
msiexec /i {product_code} /! *v "c:\removal.log" /q REMOVE="All" INSTALLATIONNAME="installation_name"
```

- To invoke a silent uninstallation and remove any queue manager data (only valid when removing the final server installation):

```
msiexec /i {product_code} /! *v "c:\removal.log" /q REMOVE="All" KEEPQMDATA="delete"
INSTALLATIONNAME="installation_name"
```

- To monitor the progress of the uninstalling process and not remove any queue manager data:

```
msiexec /x {product_code} /! *v "c:\removal.log" INSTALLATIONNAME="installation_name"
```

If you are running IBM MQ on a Windows system with User Account Control (UAC) enabled, you might see Open File - Security Warning dialog boxes during uninstallation that list International Business Machines Limited as the publisher. Click **Run** to allow the uninstallation to continue.

- To invoke a silent uninstallation and not remove any queue manager data:

```
msiexec /x {product_code} /! *v "c:\removal.log" /q INSTALLATIONNAME="installation_name"
```

## Results

After the command is entered, the command prompt immediately reappears and IBM MQ is uninstalled as a background process. If you entered parameters to produce a log, check this file to see how the uninstallation is progressing. If the uninstallation finishes successfully, you see the message Removal completed successfully in the log file.

## What to do next

Complete the steps that you started in “Uninstalling IBM MQ on Windows systems” on page 395.

## Uninstalling IBM MQ by using MQParms

You can uninstall IBM MQ by running the **MQParms** command from the command line to remove all currently installed features.

### Before you begin

Start the uninstalling process by following the steps described in “Uninstalling IBM MQ on Windows systems” on page 395.

### Procedure

1. Follow the instructions on the MQParms installation pages to uninstall IBM MQ non-interactively. See: “Using the MQParms command” on page 287.
  - a. Set the ADDLOCAL parameter to empty (ADDLOCAL="").
  - b. Set the REMOVE parameter to "ALL" (REMOVE="ALL").
2. If you have multiple versions of IBM MQ installed on your system, specify the product code that identifies the installation you want to remove. Type the following command:

```
MQParms.exe parameter_file/i "{product_code}"
```

where

- *parameter\_file* is the file that contains the required parameter values. If this file is not in the same folder as MQParms.exe, specify the full path and file name. If you do not specify a parameter file, the default is MQParms.ini.
- *product\_code* is the value shown for MSIProdCode in the output of the following command:  

```
dspmqinst -n installation_name
```

where *installation\_name* is the name of the installation you want to remove. An example of a product code is {0730749B-080D-4A2E-B63D-85CF09AE0EF0}.

### What to do next

Complete the steps that you started in “Uninstalling IBM MQ on Windows systems” on page 395.

## Uninstalling IBM MQ on Windows by using the installation media

You can uninstall IBM MQ by using the installation media to remove all currently installed features and optionally remove existing queue managers and their data.

### Before you begin

Start the uninstalling process by following the steps described in “Uninstalling IBM MQ on Windows systems” on page 395.

### Procedure

1. Insert the IBM MQ for Windows Server DVD into the DVD drive.
2. Start the installation process.
  - If autorun is enabled, the installation process starts automatically.
  - If autorun is not enabled, double-click the **Setup** icon in the root folder of the DVD to start the installation process.

The IBM MQ Installation Launchpad window opens.

3. Click **IBM MQ Installation**.
4. Click **Launch IBM MQ Installer** and click **Next** until the IBM MQ Program Maintenance panel is displayed with a welcome message. If this panel is not displayed, IBM MQ for Windows is not currently installed.

5. Click **Maintain or upgrade an existing instance** and if there is more than one installation of IBM MQ on the system, select which installation you want to remove. Click **Next** and in the Program Maintenance panel, click **Remove**, then **Next**.
6. If you are uninstalling the last or only server, and there are any queue managers on the system, the Removing Server feature panel is shown. Click one of the following options:
  - **Keep**: keep existing queue managers and their objects.
  - **Remove**: remove existing queue managers and their objects.
 Click **Next**. The Remove IBM MQ panel is displayed, with a summary of the installation to be removed.
7. Click **Remove** to continue. If there are any messages that state that locked files are found, ensure that there are no IBM MQ programs running; see “Uninstalling IBM MQ on Windows systems” on page 395. When IBM MQ has been uninstalled, a message indicates completion.
8. Click **Finish**.

## What to do next

Complete the steps that you started in “Uninstalling IBM MQ on Windows systems” on page 395.

## Uninstalling IBM MQ for IBM i



There are two ways of uninstalling IBM MQ for IBM i.

To uninstall IBM MQ for IBM i, perform one of the following tasks:

- A *standard* deletion removes IBM MQ product code but preserves user data.
- An *entire* deletion removes both IBM MQ product code and user data.

Both types of deletion require you to be signed on to the system with a user profile that has \*ALLOBJ special authority, for example, QSECOFR. Security administrator (\*SECADM) special authority is also required to delete the QMQM and QMQMADM user profiles.

### Related concepts:

“Reinstalling IBM MQ for IBM i” on page 404

You can reinstall IBM MQ for IBM i without losing any of your data.

### Related tasks:

“Uninstalling IBM MQ”

“Uninstalling IBM MQ and data” on page 402

“Uninstalling IBM MQ Java Messaging and Web Services” on page 403

Follow these instructions to uninstall IBM MQ Java.

“Uninstalling IBM MQ MQI client for IBM i” on page 403

If the IBM MQ MQI client for IBM i must be uninstalled, follow the correct procedure to ensure that all the relevant directories and files are removed.

## Uninstalling IBM MQ

### About this task

Perform a standard deletion of the IBM MQ for IBM i product if you want to retain your user data, for example, because you intend to reinstall the product at a later date.

To perform this deletion:



## Procedure

1. Quiesce IBM MQ for IBM i. (See Quiescing IBM MQ for IBM i .)
2. End the IBM MQ subsystem, by issuing the command:  
ENDSBS SBS(QMQM)
3. Ensure that no locks are held on the library QMQM, by issuing the command:  
WRKOBJLCK OBJ(QMQM) OBJTYPE(\*LIB)
4. Use the Delete Licensed Program (DLTLICPGM) command to delete the base product (and also the samples if you chose to install them).  
To delete only the samples, issue the command:  
DLTLICPGM LICPGM( 5724H72 ) OPTION(1)  
To delete only extra language versions installed, issue the command:  
DLTLICPGM LICPGM( 5724H72 ) LNG(nnnn)

where nnnn is the language number, as in the list here:

Table 64. Globalizations of IBM MQ for IBM i.

| Language ID | Language                                           |
|-------------|----------------------------------------------------|
| 2909        | Belgian English                                    |
| 2966        | Belgian French MNCS (Multi-National Character Set) |
| 2981        | Canadian French MNCS                               |
| 2975        | Czech                                              |
| 2950        | English uppercase                                  |
| 2924        | English uppercase and lowercase                    |
| 2984        | English US DBCS                                    |
| 2938        | English US uppercase DBCS                          |
| 2928        | French                                             |
| 2940        | French MNCS                                        |
| 2929        | German                                             |
| 2939        | German MNCS                                        |
| 2976        | Hungarian                                          |
| 2932        | Italian                                            |
| 2942        | Italian MNCS                                       |
| 2962        | Japanese                                           |
| 2986        | Korean                                             |
| 2978        | Polish                                             |
| 2979        | Russian                                            |
| 2989        | Simplified Chinese                                 |
| 2931        | Spanish                                            |

To delete the base product and the samples, issue the command:

```
DLTLICPGM LICPGM(5724H72) OPTION(*ALL)
```

## Results

Deleting IBM MQ for IBM i in this way deletes only the objects that belong to IBM MQ: the QMQM library, the QMQM samp library, and the subdirectories that belong to IBM MQ server within the /QIBM/ProdData/mqm directory.

If that leaves no other subdirectories (for example if IBM MQ Java is installed it uses subdirectories there) then the /QIBM/ProdData/mqm directory itself is deleted.

None of the queue manager journal libraries, or IFS directories based upon /QIBM/UserData are removed.

## Uninstalling IBM MQ and data

### About this task

You can delete IBM MQ entirely, including all user data. If you do this, save your user data first. It cannot be recovered.

To perform this deletion:

### Procedure

1. Quiesce IBM MQ for IBM i. (See Quiescing IBM MQ for IBM i .)
2. Delete each queue manager in turn by using the command WRKMQM and selecting option 4.
3. End the IBM MQ subsystem, by issuing the command:  
ENDSBS SBS(QMQM)
4. Ensure that no locks are held on the library QMQM, by issuing the command:  
WRKOBJLCK OBJ(QMQM) OBJTYPE(\*LIB)
5. Optional: If you want to also uninstall IBM MQ Java, you can do it now, using the command:  
DLTLICPGM LICPGM( 5724L26 ) OPTION(\*ALL)

This will also uninstall the Java Samples, if they were installed.

6. Use the Delete Licensed Program (DLTLICPGM) command to delete the base product (and also the samples if you chose to install them). To delete the base product and the samples issue the command:  
DLTLICPGM LICPGM( 5724H72 ) OPTION(\*ALL)
7. Delete the directory /QIBM/UserData/mqm and its subdirectories. Do this using the EDTF command and selecting option 9 (recursive delete) for the mqm directory, as follows,

**Note:** If you do this, you no longer have any information regarding your installation. Use this command with extreme caution.

The format of the command is:

```
EDTF STMF('/QIBM/UserData')
```

Alternatively, you can delete the /QIBM/UserData/mqm directory and its subdirectories by repeated use of the RMVLNK and RMVDIR commands.

8. Identify all the users who belong to the QMQMADM group. Use the DSPUSRPRF command to display a list of them. You must remove the QMQMADM group profile from their user profiles before you can delete the QMQMADM user profile. The format of the command is:  
DSPUSRPRF USRPRF(QMQMADM) TYPE(\*GRPMBR)
9. You must alter the ownership or delete the objects. For each of the user profiles QMQM and QMQMADM, use the WRKOBJOWN command to list all the objects owned by the profile. The format of the command is:  
WRKOBJOWN USRPRF( PROFILE )
10. Delete the two user profiles. The format of the command is:  
DLTUSRPRF USRPRF(QMQM) OWNNOBJOPT(\*DLT)  
DLTUSRPRF USRPRF(QMQMADM) OWNNOBJOPT(\*DLT)

## Uninstalling IBM MQ Java Messaging and Web Services

Follow these instructions to uninstall IBM MQ Java.

### About this task

To uninstall the IBM MQ Java product.

### Procedure

1. Make sure you are signed on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.
2. Issue the command:  
`DLTLICPGM LICPGM(5724L26) OPTION(*ALL)`

### Results

Deleting IBM MQ Java for IBM i deletes the objects that belong to it: the QMQMJAVA library, and the subdirectories that belong to IBM MQ Java within the /QIBM/ProdData/mqm directory.

If that leaves no other subdirectories (for example if the IBM MQ Server is installed it uses subdirectories there) then the /QIBM/ProdData/mqm directory itself is deleted.

## Uninstalling IBM MQ MQI client for IBM i

If the IBM MQ MQI client for IBM i must be uninstalled, follow the correct procedure to ensure that all the relevant directories and files are removed.

### Procedure

1. Make sure you are signed on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.
2. Use the Delete Licensed Program ( `DLTLICPGM` ) command to delete the IBM MQ MQI client for IBM i product (and also the samples if you chose to install them):  
To delete only the samples, issue the command  
`DLTLICPGM LICPGM(5725A49) OPTION(1)`  
To delete IBM MQ MQI client and the samples, issue the command:  
`DLTLICPGM LICPGM(5725A49) OPTION(*ALL)`

### Results

Deleting IBM MQ MQI client for IBM i deletes the objects that belong to it - the QMQM library, and the subdirectories that belong to IBM MQ MQI client for IBM i within the /QIBM/ProdData/mqm directory. If that leaves no other subdirectories (for example if the IBM MQ Java Client for IBM i is installed it uses subdirectories there) then the /QIBM/ProdData/mqm directory itself is deleted.

## Uninstalling IBM MQ Managed File Transfer on IBM i

Follow these instructions to uninstall IBM MQ Managed File Transfer on IBM i.

### Before you begin

To uninstall IBM MQ Managed File Transfer for IBM i, perform one of the following tasks:

- A *standard* deletion removes Managed File Transfer product code but preserves user data.
- An *entire* deletion removes both Managed File Transfer product code and user data.

Note that an entire deletion requires that you manually remove the configuration data in the /QIBM/UserData/mqm/mqft directory.

Both types of deletion require you to be signed on to the system with a user profile that has \*ALLOBJ special authority, for example, QSECOFR.

### About this task

To uninstall the IBM MQ Managed File Transfer product.

### Procedure

1. Make sure you are signed on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.

2. Issue the command:

```
DLTLICPGM LICPGM(5725M50) OPTION(*ALL)
```

### Results

Deleting IBM MQ Managed File Transfer for IBM i deletes the objects that belong to it: the QMQMMFT library, and the subdirectories that belong to IBM MQ Managed File Transfer within the /QIBM/ProdData/mqm directory.

Note that licence files are copied to /QIBM/ProdData/mqm/properties/version, and an uninstall will delete files in this directory. However, files are left in /QIBM/ProdData/mqm/properties/5725M50 as trash. For a clean uninstall, you must delete the files in this directory.

## Reinstalling IBM MQ for IBM i



You can reinstall IBM MQ for IBM i without losing any of your data.

When you reinstall IBM MQ for IBM i, the system checks whether the IBM MQ configuration file (mqc.ini) exists. If the file exists, it is kept and used with the newly installed system. If the file does not exist, an empty mqc.ini file is placed in the directory /QIBM/UserData/mqm.

All data that you have in the UserData directory is referenced by the newly installed system. In addition, all the queue manager-associated libraries containing journal and receiver information are referenced by the new system.

### Related tasks:

“Installing IBM MQ server on IBM i” on page 303

Install IBM MQ for IBM i by installing the IBM MQ server in its primary language, installing samples and installing additional languages.

---

## Uninstalling IBM MQ Telemetry components

Remove IBM MQ Telemetry components from your computer.

### Before you begin

Check that your user ID has the following authority to complete uninstallation tasks:

- On all Windows operating systems and editions, your user ID must be a member of the Administrators group.
- On AIX and Linux systems, your user ID must have root authority to complete installation. Follow your local security guidelines to acquire root authority; either login as root, or log in as another user and become root.

### Procedure

- On Windows, to uninstall IBM MQ Telemetry you must modify the existing IBM MQ. Deselect IBM MQ Telemetry from the list of components that is installed in the main installation.
- On AIX, run the command **installp -u mqm.xr.service**. For more information, refer to “Uninstalling IBM MQ on AIX” on page 389, but replace step 4 with the command in this step.
- On Linux, run the command **rpm -e MQSeriesXRService**. For more information refer to “Uninstalling IBM MQ on Linux” on page 392, but replace step 4 with the command in this step.

---

## Uninstalling IBM MQ Advanced Message Security

Information provided guides you through the uninstallation process of IBM MQ Advanced Message Security component.

### Related tasks:

“Uninstalling on AIX”

On AIX platforms, you can remove IBM MQ Advanced Message Security component either using SMIT or the command line.

“Uninstalling on HP-UX” on page 407

Use the swremove command to remove IBM MQ Advanced Message Security component on HP-UX platforms.

“Uninstalling on Linux” on page 407

Use the rpm command to remove IBM MQ Advanced Message Security component on Linux platforms.

“Uninstalling on Windows” on page 408

You can uninstall IBM MQ Advanced Message Security component using the GUI uninstallation wizard, or a command-line interface.

### Uninstalling on AIX

On AIX platforms, you can remove IBM MQ Advanced Message Security component either using SMIT or the command line.

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group mqm.

- b. Set up your environment to work with the installation you want to uninstall. Enter the following command:
 

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `. MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:
 

```
dspmqr -o installation
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:
 

```
endmqm QMgrName
```
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:
 

```
endmqm -m QMgrName
```
3. Log in as root.
  4. Uninstall IBM MQ AMS component using either **installp** or **smit**. If IBM MQ AMS component was installed in a non-default location, you must use **installp** to uninstall.
    - Uninstall using **installp** by entering one of the following commands:
      - For an installation in the default location `/usr/mqm`

```
installp -u mqm.ams.rte
```
      - For an installation in a non-default location:
 

```
installp -R usil -u mqm.ams.rte
```

where `usil` is the path of the User Specified Install Location (USIL) specified when the product was installed.
    - Uninstall using **smit**:
      - a. Select the required **smit** window using the following sequence:
 

```
Software Installation and Maintenance
Software Maintenance and Utilities
Remove Installed Software
```
      - b. List the software in the **SOFTWARE name** field:
        - 1) Enter `.`
        - 2) Press **F4**
      - c. Select the file sets to uninstall from the list (those beginning with `mqm`), and press **Enter**. There is an option at this stage to do a preview. Leave the option set to the default value of **Yes** to preview the file sets you are uninstalling, or select **No** to not preview these file sets.
      - d. Press **Enter** on the **Remove Installed Software** panel, it asks whether you are sure, press **Enter**.

## Results

The IBM MQ Advanced Message Security component has been uninstalled.

## Uninstalling on HP-UX

Use the `swremove` command to remove IBM MQ Advanced Message Security component on HP-UX platforms.

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group `mqm`.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `. MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:

```
dspmqr -o installation
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:

```
endmqm QMgrName
```
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:

```
endmqrsr -m QMgrName
```
3. Log on as root.
4. Run the following command:

```
swremove MQSERIES.MQM-AMS
```

### Results

The IBM MQ Advanced Message Security component has been uninstalled.

## Uninstalling on Linux

Use the `rpm` command to remove IBM MQ Advanced Message Security component on Linux platforms.

### Procedure

1. Stop all IBM MQ applications associated with the installation you are uninstalling.
2. For a server installation, end any IBM MQ activity associated with the installation you are uninstalling:
  - a. Log in as a user in the group `mqm`.
  - b. Set up your environment to work with the installation you want to uninstall. Enter the following command:

```
. MQ_INSTALLATION_PATH/bin/setmqenv
```

where `. MQ_INSTALLATION_PATH` refers to the location where IBM MQ is installed.
  - c. Display the state of all queue managers on the system. Enter the following command:

```
dspmqr -o installation
```
  - d. Stop all running queue managers associated with the installation you want to uninstall. Enter the following command for each queue manager:

```
endmqm QMgrName
```
  - e. Stop any listeners associated with the queue managers. Enter the following command for each queue manager:

```
endmq1sr -m QMgrName
```

3. Log in as root.
4. Run the following command:

```
rpm -e package_name
```

where *package\_name* is MQSeriesAMS-V.R.M-F

- V** Represents the version of the product that you are uninstalling
- R** Represents the release of the product that you are uninstalling
- M** Represents the modification of the product that you are uninstalling
- F** Represents the fix pack level of the product that you are uninstalling

## Results

The IBM MQ Advanced Message Security component has been uninstalled.

## Uninstalling on Windows

You can uninstall IBM MQ Advanced Message Security component using the GUI uninstallation wizard, or a command-line interface.

### Using the installation wizard Procedure

1. Insert the IBM MQ Server DVD into the DVD-ROM drive.
2. If autorun is enabled, the installation process starts.  
Otherwise, double-click the **Setup** icon in the root folder of the DVD to start the installation process.  
The IBM MQ Installation Launchpad window is displayed.
3. Click the **IBM MQ Installation**.
4. Click **Launch IBM MQ Installer**. Click **Next** until the IBM MQ Program Maintenance panel is displayed with a welcome message.  
If this panel is not displayed, IBM MQ for Windows, Version 7.5 is not installed on this machine.  
When presented with the option, select to remove/maintain or upgrade.
5. Select **Maintain or upgrade an existing instance**, then click **Next**.
6. If there are any existing queue managers, the Removing Server feature panel is displayed.  
Click one of the following options, then click **Next**:
  - **Keep** - keep existing queue managers and their objects.
  - **Remove** - remove existing queue managers and their objects.  
The Program Maintenance panel is displayed, with a summary of the installation to be removed.
7. Click **Modify** and click **Next**.
8. On the list of available IBM MQ features, click Advanced Message Security, select **Do not install this feature (remove if already intalled)**, and click **Next**. The Ready to modify IBM MQ panel appears with the summary of your changes.
9. Click **Modify** and **Next** on the following panel to continue.

## Results

Selected features of IBM MQ Advanced Message Security component have been removed.



---

## Installing IBM MQ for z/OS

Use this topic to install the IBM MQ for z/OS product on your system.

IBM MQ for z/OS uses the standard z/OS installation procedure. It is supplied with a Program Directory that contains specific instructions for installing the program on a z/OS system. You must follow the instructions in the *IBM MQ for z/OS Program Directory*. They include not only details of the installation process, but also information about the prerequisite products and their service or maintenance levels.

SMP/E, used for installation on the z/OS platform, validates the service levels and prerequisite and corequisite products, and maintains the SMP/E history records to record the installation of IBM MQ for z/OS. It loads the IBM MQ for z/OS libraries and checks that the loads have been successful. You then have to customize the product to your own requirements.

Before you install and customize IBM MQ for z/OS, you must decide the following:

- Whether you are going to install one of the optional national language features. See, National language support.
- Which communications protocol and distributed queuing facility you are going to use. See, Communications protocol and distributed queuing.
- What your naming convention for IBM MQ objects will be. See, Naming conventions.
- What command prefix string (CPF) you are going to use for each queue manager. See, Using command prefix strings.

You also need to plan how much storage you require in your z/OS system to accommodate IBM MQ ; Planning your storage and performance requirements on z/OS helps you plan the amount of storage required.

### National language support

You can choose one of the following national languages for the IBM MQ operator messages and the IBM MQ operations and control panels (including the character sets used). Each language is identified by one of the following language letters:

|   |                           |
|---|---------------------------|
| C | Simplified Chinese        |
| E | U.S. English (mixed case) |
| F | French                    |
| K | Japanese                  |
| U | U.S. English (uppercase)  |

The samples, IBM MQ commands, and utility control statements are available only in mixed case U.S. English.

### Communications protocol and distributed queuing

The distributed queuing facility provided with the base product feature of IBM MQ can either use APPC (LU 6.2), TCP/IP from IBM, or any TCP product which supports the z/OS Unix Sockets API. The distributed queuing facility is also known as the channel initiator and the mover.

You must perform the following tasks to enable distributed queuing:

- Choose which communications interface to use. This can be either, or both, of the following:

- APPC (LU 6.2)
- TCP/IP
- Customize the distributed queuing facility and define the IBM MQ objects required.
- Define access security.
- Set up your communications. This includes setting up your TCPIP.DATA data set if you are using TCP/IP, LU names, and side information if you are using APPC. This is described in Setting up communication for z/OS .

## Naming conventions

It is advisable to establish a set of naming conventions when planning your IBM MQ systems. The names you choose will probably be used on different platforms, so you should follow the convention for IBM MQ, not for the particular platform.

IBM MQ allows both uppercase and lowercase letters in names, and the names are case sensitive. However, some z/OS consoles fold names to uppercase, so do not use lowercase letters for names unless you are sure that this will not happen.

You can also use numeric characters and the period (.), forward slash (/), underscore (\_) and percent (%) characters. The percent sign is a special character to Security Server (previously known as RACF® ), so do not use it in names if you are using Security Server as your External Security Manager. Do not use leading or trailing underscore characters if you are planning to use the Operations and Control panels.

See Rules for naming IBM MQ objects for further information.

### Choosing names for queue managers and queue-sharing groups

Each queue manager and queue-sharing group within a network must have a unique name. Do not use the same name for a queue manager and a queue-sharing group. On z/OS the names of queue managers and queue-sharing groups can be up to four characters long. Each Db2 system and data-sharing group within the network must also have a unique name.

The names of queue manager and queue-sharing groups can use only uppercase alphabetic characters, numeric characters, and dollar sign (\$), number sign (#) or at sign (@); they must not start with a numeric character. Queue-sharing group names that are less than four characters long are padded internally with at signs, so do not use names ending in the at sign.

The queue manager name is the same as the z/OS subsystem name. You might identify each subsystem as a queue manager by giving it the name QM *xx* (where *xx* is a unique identifier), or you might choose a naming convention like ADDX, where A signifies the geographic area, DD signifies the company division, and X is a unique identifier.

You might want to use your naming convention to distinguish between queue managers and queue-sharing groups. For example, you might identify each queue-sharing group by giving it the name QG *xx* (where *xx* is the unique identifier).

### Choosing names for objects

Queues, processes, name lists, and clusters can have names up to 48 characters long. Channels can have names up to 20 characters long and storage classes can have names up to 8 characters long.

If possible, choose meaningful names within any constraints of your local conventions. Any structure or hierarchy within names is ignored by IBM MQ, however, hierarchical names can be useful for system management. You can also specify a description of the object when you define it to give more information about its purpose.

Each object must have a unique name within its object type. However, each object type has a separate namespace, so you can define objects of different types with the same name. For

example, if a queue has an associated process definition, it is a good idea to give the queue and the process the same name. It is also a good idea to give a transmission queue the same name as its destination queue manager.

You could also use the naming convention to identify whether the object definition is private or a global. For example, you could call a namelist `project_group.global` to indicate that the definition is stored on the shared repository.

### Application queues

Choosing names that describe the function of each queue helps you to manage these queues more easily. For example, you might call a queue for inquiries about the company payroll `payroll_inquiry`. The reply-to queue for responses to the inquiries might be called `payroll_inquiry_reply`.

You can use a prefix to group related queues. This means that you can specify groups of queues for administration tasks like managing security and using the dead-letter queue handler. For example, all the queues that belong to the payroll application might be prefixed by `payroll_`. You can then define a single security profile to protect all queues with names beginning with this prefix.

You can also use your naming convention to indicate that a queue is a shared queue. For example, if the payroll inquiry queue was a shared queue, you might call it `payroll_inquiry.shared`.

### Storage classes and coupling facility structures

The character set you can use when naming storage classes and coupling facility structures is limited to uppercase alphabetic and numeric characters. You should be systematic when choosing names for these objects.

Storage class names can be up to 8 characters long, and must begin with an alphabetic character. You will probably not define many storage classes, so a simple name is sufficient. For example, a storage class for IMS bridge queues could be called `IMS`.

Coupling facility structure names can be up to 12 characters long, and must begin with an alphabetic character. You could use the name to indicate something about the shared queues associated with the coupling facility structure (that they all belong to one suite of applications for example). Remember that in the coupling facility, the structure names are the IBM MQ name prefixed by the name of the queue-sharing group (padded to four characters with @ symbols).

### Choosing names for channels

To help you manage channels, it is a good idea if the channel name includes the names of the source and target queue managers. For example, a channel transmitting messages from a queue manager called `QM27` to a queue manager called `QM11` might be called `QM27/QM11`.

If your network supports both TCP and SNA, you might also want to include the transport type in the channel name, for example `QM27/QM11_TCP`. You could also indicate whether the channel is a shared channel, for example `QM27/QM11_TCP.shared`.

Remember that channel names cannot be longer than 20 characters. If you are communicating with a queue manager on a different platform, where the name of the queue manager might contain more than 4 characters, you might not be able to include the whole name in the name of the channel.

### Using command prefix strings

Each instance of IBM MQ that you install must have its own *command prefix* string (CPF). You use the CPF to identify the z/OS subsystem that commands are intended for. It also identifies the z/OS subsystem from which messages sent to the console originate.

You can issue all MQSC commands from an authorized console by inserting the CPF before the command. If you enter commands through the system command input queue (for example, using CSQUTIL), or use the IBM MQ operations and control panels, you do not use the CPF.


To start a subsystem called CSQ1 with CPF that is '+CSQ1 ', issue the command +CSQ1 START QMGR from the operator console (the space between the CPF and the command is optional).

The CPF also identifies the subsystem that is returning operator messages. The following example shows +CSQ1 as the CPF between the message number and the message text.

```
CSQ9022I +CSQ1 CSQNCDSP ' DISPLAY CMDSERV' NORMAL COMPLETION
```

See the Defining command prefix strings (CPFs) for information about defining command prefix strings.

**Related information:**

 [Program Directory for IBM MQ for z/OS](#)

---

## Planning to install IBM MQ

To install the IBM MQ product your hardware, and software environment must meet minimum requirement levels. You must also consider the national language features, communications protocols, and naming conventions to be used.

### IBM MQ Prerequisites

Before attempting to install, and run IBM MQ for z/OS, ensure that your system hardware, and software levels meet the minimum requirement. You can check the minimum required levels at IBM MQ for z/OS requirements.

Note that the System Requirements page for IBM MQ Version 8.0 uses the Software Product Compatibility Reports (SPCR) tool.

The SPCR tool enables you to go directly to the:

- Supported operating systems
- Prerequisites
- System requirements, and
- Optional supported software.

---

## Delivery media

IBM MQ for z/OS is supplied on 3590 cartridge only.

One cartridge contains the product code together with the following language features; U.S. English (mixed case), U.S. English (uppercase), French, Chinese, and Japanese.

---

## Customizing IBM MQ and its adapters

IBM MQ requires some customization after installation to meet the individual and special requirements of your system, and to use your system resources in the most effective way.

For a list of tasks that you must perform when you customize your system, see Customizing IBM MQ for z/OS.

## Using queue-sharing groups

If you want to use queue-sharing groups, you do not have to set them up when you install IBM MQ, you can do this at any time.

For details of how to manage your queue-sharing groups when you have set them up, see Managing queue-sharing groups.

---

## Verifying your installation of IBM MQ for z/OS

After the installation and customization has been completed, you can use the installation verification programs (IVPs) supplied with IBM MQ for z/OS to verify that the installation has been completed successfully.

The IVPs supplied are assembler language programs and you should run them after you have customized IBM MQ for z/OS to suit your needs. They are described in Running the basic installation verification program.

---

## Macros intended for customer use

The macros identified in this topic are provided as programming interfaces for customers in support of features that are specific to IBM MQ for z/OS.

The 'C' include files, COBOL copy files, PL/I include files and assembler macros that are provided as programming interfaces for customers in support of features that apply across many IBM MQ platforms are described in the Constants.

**Note:** Do not use as programming interfaces any IBM MQ macros other than those interfaces identified in this topic or in the Constants

## General-use programming interface macros

The following assembler macros are provided to enable you to write programs that use the services of IBM MQ. The macros are supplied in library thlqual.SCSQMACS.

- CMQXCALA
- CMQXCFBA
- CMQXCFCFA
- CMQXCFLA
- CMQXCDFFA

- CMQXCINA
- CMQXCVCA

## Product-sensitive programming interface macros

The following assembler macros are provided to enable you to write programs that use the services of IBM MQ. The macros are supplied in library thlqual.SCSQMACS. Product-sensitive interfaces are open to change between different releases of the product.

- CSQBDEF
- CSQDQEST
- CSQDQIST
- CSQDQJST
- CSQDQLST
- CSQDQMAC
- CSQDQMST
- CSQDQPST
- CSQDQSST
- CSQDQWHC
- CSQDQWHS
- CSQDQ5ST
- CSQDWQ
- CSQDWTAS
- CSQQDEFX
- CSQQLITX

---

## Sub-capacity license charges with IBM MQ for z/OS

Sub-capacity license charges is a particular way of charging you for an IBM product that runs on a z/OS system, which is based on how much use you make of the product.

To determine the product usage, the z/OS system records the amount of processor time that is used by the product when it executes. z/OS can measure how much processing time is spent in doing work on behalf of the IBM MQ queue manager that is handling MQI calls, executing MQSC commands, or performing some other action to support the messaging and queuing functions that are used by your application programs. The amount of processing time is recorded in a file at hourly intervals, and the hourly records are totaled at the end of a month. In this way, the total amount of time used by the IBM MQ for z/OS product on your behalf is computed, and used to determine how much you pay for your use of the IBM MQ for z/OS product that month.

Sub-capacity license charges are implemented as follows:

- When IBM MQ for z/OS is installed, it identifies itself to z/OS and requests that the *System Management Facilities (SMF)* mechanism within z/OS is to automatically measure how much processor time is used by the IBM MQ for z/OS product.
- When enabled, the z/OS usage measurement facility collects usage figures for each hour of the day, and generates usage records that are added to a report file on disk.
- At the end of one full month, these usage records are collected by a program, which generates a report of product usage for the month. This report is used to determine the charge for the IBM MQ for z/OS product.

**Note:** A VUE-enabled queue manager records usage information in SMF89 records with the product name and identifier for IBM MQ for z/OS Value Unit Edition (VUE) instead of those for the IBM MQ product.

For more information about sub-capacity license charging and the Sub-Capacity Reporting Tool (SCRT), see IBM Z Software Pricing. For information about the MULCCAPT parameter see, Using CSQ6SYSP.

---

## IBM MQ for z/OS Value Unit Edition (VUE)

### IBM MQ for z/OS Value Unit Edition (VUE)

IBM MQ for z/OS Value Unit Edition (VUE) provides all the function and capability of the base IBM MQ for z/OS, in a format that offers a one-time-charge (OTC) price metric for eligible workloads that are deployed in qualified IBM Z Systems New Application License Charge (zNALC) logical partitions (LPARs).

The term, Eligible Workload, is defined as new workload that executes using the IBM MQ for z/OS VUE server environment, on condition that the workload is qualified and approved through the zNALC qualification process.

For further information on zNALC, see IBM Z Systems Software Pricing

The OTC price metric provides an alternative pricing model for new IBM MQ for z/OS-connected applications and new IBM MQ for z/OS VUE service enablement workloads.

Support for the zNALC metric offers a reduced price for the z/OS operating system on LPARs that run a qualified application.

IBM MQ for z/OS VUE can connect to other supported versions of IBM MQ for z/OS (whether in zNALC or non-zNALC environments) for workload federation and systems management.

IBM MQ for z/OS VUE allows connections from IBM MQ clients, that run on other platforms.

### Installing VUE

An order for VUE is fulfilled by delivery of two products:

- IBM MQ for z/OS,
- The VUE enabling product itself.

The products are separately installed using SMP/E following the process documented in their respective program directories.

### Enabling VUE

To enable a queue manager to run an eligible workload in a zNALC LPAR, the SCUEAUTH library created by installation of the VUE enabling product must be added to the STEPLIB concatenation of the xxxxMSTR procedure for that queue manager:

- The SCUEAUTH library should be APF authorized
- The SCUEAUTH library must be concatenated ahead of the SCSQAUTH library,

For example, the CSQ4MSTR sample would be modified as follows:

```
//PROCSTEP EXEC PGM=CSQYASCP,REGION=0M,MEMLIMIT=2G
//*
//STEPLIB DD DSN=h1q.SCSQANLE,DISP=SHR
// DD DSN=h1q.SCUEAUTH,DISP=SHR
// DD DSN=h1q.SCSQAUTH,DISP=SHR
-
```

When SCUEAUTH is added to the STEPLIB concatenation, the characters VUE appear in the release level shown by message CSQY000I during queue manager startup. The queue manager starts only in an LPAR configured for zNALC workload by name or IPL parameter.

## Characteristics of a VUE-enabled queue manager

A VUE-enabled queue manager has all the function and capability of the base queue manager. Additionally, clients will be enabled during channel initiator start up.

A VUE-enabled queue manager records usage information in SMF89 records with the product name and identifier for IBM MQ for z/OS Value Unit Edition (VUE) instead of those for the IBM MQ product.

A VUE-enabled queue manager can:

- Connect to other queue managers and clients in a network, according to the connectivity capabilities of the base queue manager installation.
- Participate in a queue sharing group with other queue managers provided the base queue manager versions are able to interoperate, regardless of whether other members are standard or VUE function queue managers.

---

## IBM MQ Advanced Message Security for z/OS

IBM MQ Advanced Message Security for z/OS (AMS) is a separately licensed enabling product that extends IBM MQ to provide a high level of protection for sensitive data flowing through the IBM MQ network using a public key cryptography model.

### Installing AMS

IBM MQ Advanced Message Security for z/OS is installed separately using SMP/E by following the process documented in its program directory.

### Enabling AMS

IBM MQ Advanced Message Security for z/OS is enabled separately for each queue manager by completing additional customization tasks described at Customizing IBM MQ for z/OS.

The following tasks are relevant when adding AMS support to a queue manager:

- Task 2: APF authorize the IBM MQ load libraries
- Task 3: Update the z/OS link list and LPA
- Task 4: Update the z/OS program properties table
- Task 13: Customize the initialization input data sets
- Task 17: Tailor your system parameter module
  - Using CSQ6SYSP
- Task 23: Create procedures for Advanced Message Security
- Task 24: Set up the started task user Advanced Message Security
- Task 25: Grant RACDCERT permissions to the security administrator for Advanced Message Security
- Task 26: Grant users resource permissions for WebSphereMQ Advanced Message Security




You also need to configure certificates and policies, which are described in

- Using certificates on z/OS
- Security policies
- Example configurations on z/OS

**Related information:**

IBM MQ Advanced Message Security

 [IBM MQ Advanced Message Security for z/OS V8.0 Program Directory \(GI13-3329-00\)](#)

---

## IBM MQ Managed File Transfer for z/OS

Use this topic to install IBM MQ Managed File Transfer on your IBM MQ for z/OS system.

### Overview

IBM MQ Managed File Transfer for z/OS uses the standard z/OS installation procedure. The *Program Directory for IBM MQ Managed File Transfer for z/OS* contains specific instructions for installing the program. You should follow the instructions in this document. The instructions include not only details of the installation process, but also information about the prerequisite products and their service or maintenance levels.

SMP/E, used for installation on the z/OS platform, validates the service levels and prerequisite and corequisite products, and maintains the SMP/E history records to record the installation of IBM MQ Managed File Transfer. The process loads the appropriate libraries and checks that the loads have been successful. You then have to customize the product to your own requirements.

**Note:** For Version 8.0, the supported version of Java for IBM MQ Managed File Transfer for z/OS is Java 1.7.

### Planning

See Planning for IBM MQ Managed File Transfer for items you need to consider before installing the component.

### Installing

Install the product by following the instructions detailed in the *Program Directory for IBM MQ Managed File Transfer for z/OS*.

Check that the SMP/E installation process has created the product JCL library USERID.MFTV800.SBFGCMDS.


If this JCL library has not been created during the installation process, create the library and submit the job USERID.ZOS.JCL(COPYJCL1).

When you have installed the product, you must carry out some customization tasks. See IBM MQ Managed File Transfer for z/OS for more information.

**Related information:**

IBM MQ Managed File Transfer for z/OS

Planning for IBM MQ Managed File Transfer

 [IBM MQ Managed File Transfer for z/OS V8.0 Program Directory \(GI13-3330-00\)](#)


---

## Migrating and upgrading IBM MQ

Migration is the process of updating queue managers, and other objects, such as applications or administrative procedures. To migrate a queue manager to run on a new level of code, you must first upgrade IBM MQ to install the new code level. When you have verified that the upgrade is successful, migrate the queue manager and all the applications and resources that are associated with it. Before you start this process, create a migration plan, based on the information in this documentation.


### Maintenance, migration and upgrading


IBM MQ uses this terminology as follows:


- Migration is the process of updating queue manager data to match a newer level of code. This occurs the first time a queue manager is started with the newer level of code.
- Maintenance is the application of a fix pack, interim fix or program Temporary Fix (PTF). The installation can be restored to its previous level and queue managers or applications continue to work. Migration is not required after applying maintenance. However, you should test applications with the new level of IBM MQ code.
- Upgrading is the process of taking an existing IBM MQ installation and upgrading to a new level of code. Unless the upgrade is applying a fix (and not enabling new function), an upgrade must be followed by migration.
- Once migration has occurred, the queue manager can no longer be started by an earlier code level. On most platforms, queue manager migration is not reversible.  The exception is z/OS, where you can reverse queue manager migration, but only if you have not enabled new function.

### IBM WebSphere MQ / IBM MQ Migration Guides

 The *IBM WebSphere MQ / IBM MQ Migration Guide* provides information to help you plan the process of migrating from an older version to a new version of the product on distributed systems.

- For an introduction to the guide and its contents, see the developerWorks® blog article [IBM WebSphere MQ / IBM MQ Migration Guide](#).
- To view the guide in your web browser, click the following link: [IBM WebSphere MQ / IBM MQ Migration Guide - HTML version](#).
-  To download the guide as a PDF file, click the following link: [IBM WebSphere MQ / IBM MQ Migration Guide - PDF file](#).

 The *IBM WebSphere MQ/ IBM MQ for z/OS Migration Guide* provides information to help you plan the process of migrating from an older version to a new version of the product on z/OS.

- For an introduction to the guide and its contents, see the developerWorks blog article [IBM WebSphere MQ / IBM MQ for z/OS Migration Guide](#).
- To view the guide in your web browser, click the following link: [IBM WebSphere MQ / IBM MQ for z/OS Migration Guide - HTML version](#).
-  To download the guide as a PDF file, click the following link: [IBM WebSphere MQ / IBM MQ for z/OS Migration Guide - PDF file](#).

### Where to find information about system requirements and prerequisites

From IBM MQ Version 8.0, you can use the Software Product Compatibility Reports (SPCR) tool to find information on supported operating systems, system requirements, prerequisites, and optional supported

software. For more information about the SPCR tool and links to reports for each supported platform, see the System Requirements for IBM MQ Version 8.0 web page.

For links to system requirements information for all releases of IBM WebSphere MQ or IBM MQ, see the IBM MQ System Requirements web page.

For information about limitations and known problems for IBM MQ Version 8.0 and its maintenance, see the product readme file, which is available from the product readmes web page.

## **Migrating from a release of IBM WebSphere MQ before Version 7.0**

**Important:** If you are migrating your system from a version of IBM WebSphere MQ before Version 7.0, you must migrate your system to Version 7.0, Version 7.0.1, or Version 7.1 before you migrate to Version 8.0. See the appropriate version of the product documentation for information on how to carry out the task. For links to earlier versions of the product documentation not available in IBM Knowledge Center, see the IBM MQ documentation library page.

## **Getting started with migrating and upgrading**

If you are not familiar with IBM MQ migration, start by reading the following information:

- The “Introduction to IBM MQ migration” on page 421 section: use these topics to find out more about the concepts that you must understand before planning migration tasks, including the difference between maintenance, migration, and upgrading and which migration paths are supported.
- The “IBM WebSphere MQ / IBM MQ Migration Guides” on page 419: use these guides to find more information about planning the migration process for your release and platform.



For information about new features and changes in this release, see the following topics:

- What's new in IBM MQ Version 8.0
- What's changed in IBM MQ Version 8.0
- What's new and what's changed in MQ Explorer

Some of the changes have a migration impact, because they affect the behavior of existing applications or the automation of management tasks. This subset of the changes is listed in “Changes that affect migration” on page 580. Study this list of changes to plan what migration tasks you must perform.

If you prefer to use the IBM MQ Version 8.0 product documentation offline, you can download it through the links on the IBM MQ documentation library page, either as a downloadable package or as a set of PDF files.

## Related information:

  IBM MQ for z/OS Program Directory

## Introduction to IBM MQ migration

These topics explain the concepts that you must understand before planning migration tasks, where to find migration topics, and which migration paths are supported.

## Migration paths

You can find topics about direct migration to the current release of IBM MQ in this release of the IBM MQ product documentation. Paths between other releases are in previous releases of the product documentation, which are available on the web.

### Migration paths: IBM MQ (On platforms except z/OS )

**Note:**   Backward migration is not possible

It is impossible to restore a queue manager from Version 7.0.1 to an earlier version. If a queue manager has not been started, it is possible to uninstall the current version and reinstall a different version of IBM MQ. It does not matter what versions of IBM MQ are installed between when a queue manager was last started and when it is next started.

Table 65. Migration paths: IBM MQ (On platforms except z/OS )

| From / To              | Version 7.R                                                                                                                                                                                    | Version 8.0                                                                                                                                |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Prior to Version 7.0.1 | Follow the instructions for your current version. For more information about where to find the documentation for older versions of the product, see Documentation for older versions of IBM MQ | Direct migration is not possible. You must follow an indirect migration path, which involves migrating IBM MQ more than once.              |
| Version 7.0.1 or later | Follow the instructions for your current version. See IBM MQ product documentation in IBM Knowledge Center.                                                                                    | “Introduction to IBM MQ migration” and “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485 |



### Migration paths: IBM MQ for z/OS

Table 66. Migration paths: IBM MQ for z/OS

| From / To              | Version 7.1                                                                                                   | Version 8.0                                                                                                                                                                                                                                    |
|------------------------|---------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prior to Version 7.0.1 | Follow the instructions on the Web for your current version. See the IBM MQ library page for further details. | Direct migration is not possible. You must follow an indirect migration path, which involves migrating IBM MQ more than once.<br><br>See “Migrating from earlier unsupported releases of IBM MQ for z/OS” on page 551 for further information. |

Table 66. Migration paths: IBM MQ for z/OS (continued)

| From / To              | Version 7.1                                                                                                 | Version 8.0                                                  |
|------------------------|-------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| Version 7.0.1 or later | Follow the instructions for your current version. See IBM MQ product documentation in IBM Knowledge Center. | "z/OS: Migration planning to the latest release" on page 486 |

To revert an IBM MQ for z/OS queue manager to an earlier version, see "z/OS: Reverting a queue manager to a previous release" on page 563.

## The version naming scheme for IBM MQ for z/OS

On IBM MQ for z/OS, releases have a three-digit Version, Release, and Maintenance (VRM) level code. The code is significant; it identifies the service life of a release. To run a queue manager at a different VRM level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

The release level of IBM MQ for z/OS is described by a three-digit VRM code. The VRM acronym stands for:

*Version . Release . Modification*

6.0.0, 7.0.0, 7.0.1, 7.1.0, and 8.0.0 are examples of IBM MQ for z/OS release level codes. On z/OS, a release of IBM MQ always has a three-digit VRM code, even if the release is the first release in a version, such as 7.0.0. IBM MQ for z/OS follows a convention of changing the VRM when the product is installed by SMP/E with a new FMID.

You can modify existing libraries, without changing the FMID, by applying PTFs. You cannot upgrade existing libraries to an FMID or release level by applying PTFs.

The release level of a z/OS queue manager is written to the operator console in the message CSQY000I.

The command level of a queue manager is a three-digit VRM code. You can look at the queue manager command level in the queue manager property panel in IBM MQ Explorer. An IBM MQ program can call MQINQ, passing the MQIA\_COMMAND\_LEVEL selector, to obtain the command level of the queue manager it is connected to.

The VRM code, or release level, is significant in two respects. Changing the release level that a queue manager runs at, requires migration of the queue manager. It also requires attention to the PTF level of other queue managers that are in the same queue sharing group. It is also significant because each release level has its own service life, and end of service date.

The service life depends on the VRM. Each release level has its own service end date. So, for example, 7.0.0, on z/OS, has a different service end date from 7.0.1. See IBM Software Support lifecycle policy. Click through to WebSphere product lifecycle dates, to find the service life and end of service for different releases of IBM MQ.

## Related concepts:

“Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance.

Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

## The version naming scheme for IBM MQ (On platforms other than z/OS )

IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

The full version of IBM MQ (On platforms other than z/OS ) is described by a four-digit VRMF code.

The version and release parts of the code are significant; they identify the service life of a release. To run a queue manager at a different VR level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

The VRMF acronym stands for:

*Version . Release . Modification . Fix*

8.0, 7.0.0.1, and 7.0.1.0 are examples of full IBM MQ version codes.

You can find the full version level of an IBM MQ installation by typing the command **DSPMQVER**, or **DSPMQMVER** on IBM i. It returns the full four-digit VRMF code.



Versions and releases of IBM MQ are known by the first two digits of the VRMF code. The two digits are sometimes prefixed by a V, such as V8.0. A version of IBM MQ always has a release level, even if it is the first release in a version.

The first release is normally labeled Vx.0, for example, IBM MQ V7.0. Occasionally, the first release of a version on a specific platform is not labeled V x.0. It is numbered to correspond the command level that has been implemented on the platform.

In documentation, the release level is sometimes dropped from the VRMF code, for example, V7. Dropping the release level can lead to ambiguity, if the context is not clear. For example, V7 might mean the whole of V7, or the release level V7.0, as opposed to the release level V7.2, or V7.3.

The third digit in the VRMF identifies the modification level of a release. A change in the third digit does not change the release. For example, after upgrading IBM MQ to modification level 7.0.1, the release of IBM MQ remains 7.0. However the command level does change to 7.0.1.

The significance of the distinction between release and modification level concerns migration, and the service life of a product. Queue manager objects, such as queue managers, channels, queues, and messages do not require migration to upgrade to a new modification level. Nor do they require migration if the modification level is removed<sup>1</sup>. Migration might be required for a version or release level change.

**Note:**   Backward migration is not possible. To be able to restore an earlier version or release level of a queue manager, you must back it up before upgrading. If you do restore it, you restore the queue manager, and its data, to the state it was in when you backed it up.

---

1. Applications using new functions introduced in a modification level do not work on an earlier level.

A new version or release has a new service end date. New modification levels generally do not result in a new service end date. But if a modification level is announced, then a new service end date might be announced too.

The fourth digit in the VRMF code is the fix level. Fix levels do not affect the command level of the queue manager. No migration is required, and fix levels do not affect the service end date of a release.

Trailing zeros in the VRMF code are never significant, but are sometimes quoted for clarity. For example, you might see 7.0.0 to distinguish it from 7.0.1, and 7.0.1.0 to distinguish it from 7.0.1.1. 7.0.0 is no different from 7.0 or 7.0.0.0, and 7.0.1 and 7.0.1.0 are the same level.

Modification levels and fix levels are known by three and four-digit VRMF codes. 7.0.1 is a modification level and 7.0.1.2 is a fix level. Modification levels are shipped as refresh packs, and fix levels as fix packs.

A refresh or fix pack is named using a two part name that uniquely identifies it. The first part of the name is a truncated VRMF. The second part of the name is the name new refresh or fix pack. So, for example, the name of the fix pack 7.0.1.2 for Windows is 7.0.1-WS-MQ-Windows-FP0002, and the name of the refresh pack 7.0.1 for Windows is 7.0-WS-MQ-Windows-RP0001.

Refresh packs and fix packs for a particular version/release are cumulative, from the initial release. You can apply any higher numbered refresh, or fix pack, of the same version/release to upgrade directly to that version level. You do not have to apply the intervening fixes. Refresh packs and fix packs are obtained as service through Fix Central.

The latest modification level is also used to refresh the version of IBM MQ available through Electronic Software Download using Passport Advantage, or on physical media.

When you order IBM MQ you receive a manufacturing refresh at the latest modification level. The result of installing a manufacturing refresh is almost the same as applying the refresh pack to an earlier fix level of IBM MQ. There is one important difference. Refresh packs are applied using a maintenance procedure, manufacturing refreshes are installed using an installation procedure. You can "unapply" a refresh pack to return to the previous fix level you had installed. You can only uninstall a manufacturing refresh, which removes IBM MQ from your system.

In addition to fixes packaged as refresh packs and fix packs, you can also obtain interim fixes for IBM MQ. You get these from Fix Central. Interim fixes are also known as emergency or test fixes, and are known collectively as interim fixes. The naming scheme for refresh and fix packs extends to interim fixes. Interim fixes are known either by their fix name, or by the list of APARs they fix.

When you apply new fix packs or refresh packs, all interim fixes are removed. The documentation with the fix pack or refresh pack tells you if the APARS associated with the interim fixes you have applied have been fixed. If they have not, check to see if there are new interim fixes, at the new level, for the APARs that concern you. If there are not, consult service. They might either tell you to reapply the interim fix, or supply a new interim fix.



**Related concepts:**

“Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )” on page 453  
You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

## Internet Protocol Version 6 (IPv6) migration

This section deals with using IPv4 and IPv6 when you are thinking of installing IBM MQ

### General Introduction

The Internet Protocol Version 6 (IPv6) is designed by the Internet Engineering Task Force (IETF) to replace the current version Internet Protocol, Version 4 (IPv4). IPv4 has been around for over 20 years and is one of the primary methods for machines to communicate to each other over the internet. IPv4 is limited to 32-bit addressing for internet addresses. These addresses are needed by all new machines added to the internet and they are beginning to run out. The IETF is the controlling standards body for the Internet and to meet the growing demand for internet addresses has increased the number of digits used for Internet addresses from 32 to 128 bits. IPv6 offers a far larger number ( $2^{128}$ ) of internet addresses and should solve the address shortage for the foreseeable future. IPv6 is expected to gradually replace IPv4, with the two protocols coexisting for a number of years while this transition period exists. IPv6 also simplifies header formats and improves support for extensions and options, flow labeling capability, and consolidated authentication and privacy capabilities

IBM MQ has the ability for queue managers to communicate using the IPv6 protocol in addition to the existing, IPv4, protocol.

Further information on IPv6 can be found at [IPv6 and IPv6 Forum](#).

### IBM MQ platforms that support IPv6

This section lists the IBM MQ platforms that support IPv6.

IPv6 is supported on the following IBM MQ platforms:

- IBM MQ for AIX
- IBM MQ for Linux
- IBM MQ for Sun Solaris
- IBM MQ for HP-UX
- IBM MQ for Windows
- IBM MQ for IBM i
- IBM MQ for z/OS

## Key points in migrating to IPv6 and using IBM MQ

This section lists some key points to be aware of when you are thinking of installing IBM MQ and using IPv6.

- IBM MQ recognizes IPv6 hexadecimal addresses (for example fe80:43e4:0204:acff:fe97:2c34:fde0:3485) as well as IPv4 dotted decimal addresses (for example 9.20.9.30).
- For a system running both IPv4 and IPv6 system, the connection name (CONNNAME) you specify for a given channel determines the IP protocol for the channel making a connection.

## Considerations when implementing IPv6 in a network

This section lists some things that you should consider when you are thinking of installing IBM MQ on an IPv6 network.

- To ensure consistency across the network, you should plan the introduction of IPv6 for the whole network, especially where clusters are involved. For example, although a queue manager is now IPv6 capable, this doesn't imply that the queue managers it can communicate with are also IPv6 capable.
- When setting the domain name server (DNS) or equivalent, consider whether the system on which the target queue manager is running can resolve to an IPv4 address, an IPv6 address or a dual IPv4 and IPv6 address.
- If the system that you are installing IBM MQ on does not support IPv6, IBM MQ will only be able to connect using IPv4.
- For a queue manager running on an IPv6 enabled system to be able to communicate with a queue manager running on an IPv4 enabled system, the IPv4 enabled system must have a hostname that resolves to an IPv4 address only.
- If there are multiple domain name servers in an IBM MQ network, each hostname used in a channel definition must resolve to the same address (or addresses), regardless of which DNS is used.

## Migrating a queue manager to IPv6

This section deals with migrating a queue manager when you are thinking of installing IBM MQ on an IPv6 network.

The IPv6 protocol can only be utilized by IBM WebSphere MQ Version 6.0 or later. In order to make use of the IPv6 protocol, IBM MQ must be installed on a system that is IPv6 capable.

The preferred IP version that two systems use for communicating (if both IPv4 and IPv6 are available) is determined by a new queue manager attribute IPADDRV. This parameter only has an effect if the hostname resolves ambiguously to both an IPv4 address and an IPv6 address.

To migrate a queue manager to use the IPv6 protocol:

1. Configure dual IPv4 and IPv6 protocols on the system where the queue manager to be migrated resides.
2. Install IBM MQ.
3. Add an entry to the DNS to resolve the hostname of the system that is to be migrated, to both an IPv4 address and an IPv6 address.
4. Set the IPADDRV parameter to IPv6 (or set the LOCLADDR parameter to resolve to an IPv6 address).

### CAUTION:

**Not all IPv6 software can interpret an IPv4 mapped IPv6 address. If the combination of CONNNAME and LOCLADDR results in an IPv4 mapped IPv6 address, ensure that the system hosting the target queue manager is capable of handling this.**

**Using mapped addresses can require protocol translators in the IP network.**

## Migration scenarios (non-cluster topology)

It is possible to come up with a number of different interconnection possibilities, and the following sections aim to help you understand how IBM MQ will work in each case.

### Non-cluster migration scenario 1

Three systems exist that are IPv4 only capable. Each system hosts a queue manager (QM1, QM2, and QM3) and each queue manager connects to the other two. All CONNAMES in the cluster channel definitions are made using DNS names rather than IP addresses.

Enable QM1 to be able to use channels running over IPv6 as follows

1. Upgrade the host system to have dual IPv4 and IPv6 stacks.

**Important:** A listener is required for each IP stack.

2. Install the latest version of IBM MQ.
3. Update the DNS table so that it has two entries for the system running QM1; one entry for its IPv4 address and one for its IPv6 address. This enables a DNS name request to return both IPv4 and IPv6 addresses for this host.
4. Set the queue manager IPADDRV attribute to IPv6.

**Note:** Even with these changes made to support IPv6 addressing, QM1 will still be able to communicate with queue managers (both existing and new ones) that are only IPv4 capable.

Enable QM2 to be able to use channels running over IPv6 as for QM1 above.

- Communications between QM1 and QM2 will now be over IPv6.
- Communications between QM1 and QM3 will still be over IPv4.
- Communications between QM2 and QM3 will still be over IPv4.

With the queue manager IPADDRV attribute set to IPv6, the preference has been set for the queue manager to connect using the IPv6 protocol. If a channel from QM1 to QM3 has LOCLADDR set to a host name which resolves to an IPv6 address, or both IPv4 and IPv6 addresses (with the IPADDRV attribute set to IPv6, the IPv6 address will be returned as that is the preference), this channel will attempt to use the IPv6 protocol. If the IPv6 protocol installed on the QM1 host system is capable of using a mapped address then QM1 will communicate with QM3 over IPv6. Otherwise, the channel will fail to resolve CONNNAME.

While QM3 remains a queue manager on an earlier version of the product, you will need to check that all CONNAMES used to start a channel to QM3 do not resolve to an IPv6 address or dual IPv4 and IPv6 addresses where the IPv6 address could be returned. This would cause QM1 to attempt to start the channel over IPv6 which would fail, as it would be unable to resolve the CONNNAME.

It is possible to upgrade a system to have dual IPv4 and IPv6 capability and still run a queue manager on an earlier version of the product, on the system. While it is not recommended to run this type of configuration, as long as the addresses that are returned to this level of queue manager are either IPv4 or an IPv4 mapped version of an IPv6 address, this should work.

### Non-cluster migration scenario 2

Three systems exist that are IPv4 only capable. Each system hosts a queue manager (QM1, QM2, and QM3) and each queue manager connects to the other two. All CONNAMES in the cluster channel definitions are made using IP addresses.

Because addresses have been specified instead of DNS names, to allow a queue manager to connect to another using the IPv6 protocol you will need to duplicate the definitions that use IPv4 addresses between them and provide them with IPv6 addresses instead. The original definitions that use IPv4 addresses will continue to work, but if you intend to take advantage of the IPv6 protocol, you will need to connect using the new definitions.

Enable QM1 to be able to use channels running over IPv6 as follows

1. Upgrade the host system to have dual IPv4 and IPv6 stacks.

**Important:** A listener is required for each IP stack.

2. Install IBM MQ.
3. Duplicate the channel, transmission queue and, where applicable, any process definitions using IPv6 addresses where required.

**Note:** Even with these changes made to support IPv6 addressing, QM1 will still be able to communicate with existing queue managers that are only IPv4 capable.

Enable QM2 to be able to use channels running over IPv6 as for QM1 above.

1. Upgrade the host system to have dual IPv4 and IPv6 stacks.

**Important:** A listener is required for each IP stack.

2. Install IBM MQ.
3. Where necessary amend applications to write to the new remote queue (created above for QM1 with the IPv6 addresses).
4. Verify the channels can be started.

The queue managers can now connect as follows:

- QM1 can now connect with QM2 over either IPv4 or IPv6 depending on the channel the application writes its messages to.
- QM1 still connects with QM3 over IPv4 using the original definitions.

## Migrating a cluster to IPv6

This section deals with migrating clusters when you are thinking of installing IBM MQ on an IPv6 capable network.

The following gives an overview of approaches that can be taken when migrating a cluster to the latest version of IBM MQ. Due to the variations that can occur within a cluster, the detail is deliberately general and should only be seen as a guide to the likely course of action you will need to take.

### Migration scenarios (cluster topology)

Where an IPv6 capable system is to be added to an IBM MQ cluster, all full repository systems in that cluster must be IPv6 capable.

The following scenarios are seen as the ones most likely to occur in customer installations. They describe the changes that are likely to be required.

#### Scenario 1

A cluster from an earlier version of the product is installed on IPv4 only capable, systems and you need to connect an IPv6 only capable system into the cluster. All CONNAMES in cluster channel definitions are made using DNS names rather than IP addresses.

When adding a new IPv6 only system to the cluster, identify those queue managers that your new system will communicate with. These include:

- The queue managers your new system will send messages to.
- The queue managers your new system will receive messages from.
- The full repository queue managers

The systems that you have identified must be upgraded before introducing the new system.

Recommended migration procedure:

- Upgrade each of the systems hosting a full repository queue manager as shown in "Migrating a queue manager to IPv6" non-cluster scenario 1.
- Upgrade the remaining cluster systems which need to be IPv6 capable as shown in "Migrating a queue manager to IPv6" non-cluster scenario 1.

With this configuration:

- The new IPv6 only capable system will communicate with the cluster using IPv6 addressing
- All other IPv4 systems that connect into the cluster will continue to communicate using IPv4 addressing
- The systems in the cluster will be able to connect to each other using either IPv4 or IPv6 addressing. The decision as to which address is used depends on whether you have set IPADDRV to specify IPv4 or IPv6 connections.


## Scenario 2

A cluster from an earlier version of the product is installed on IPv4 only capable systems and you need to connect an IPv6 only capable system into the cluster. Your network does not support adding both IPv6 and IPv4 addresses using the same hostname or you are using IP addresses rather than DNS names in the cluster channel CONNAMES.

The problem here is likely to be that all of the systems cannot be switched to IPv6 simultaneously and some at least must remain only IPv4 capable. The systems that your new IPv6 only system communicates with must be IPv4 and IPv6 capable. We do not recommend simply adding a new set of IPv6 channels into the cluster for the IPv6 system to use, as the IPv4 system would also try to use them, resulting in communication errors.

The recommended approach is:

- Define a new cluster which contains the IPv6 only capable system or systems with new IPv6 addresses and channel definitions. The existing cluster remains, and contains the IPv4 only system definitions. The image below gives a pictorial representation of this. QM1, QM2, and QM3 represent the original IPv4 cluster. QM2, QM3, and QM4 represent the new cluster created to allow the IPv6 only capable system (QM4) to connect into your configuration.
- If you are using DNS names, you can give each of the systems separate DNS names for IPv4 and IPv6 (for example system1\_IPv4.ibm.com and system1\_IPv6.ibm.com).
- Define a new CLUSRCVR channel and any corresponding CLUSSDR channels using the new IPv6 names or IP addresses on each system in the new cluster. In this way the systems with only IPv4 or IPv6 capability do not see channels which they are not able to use and no communications error will result.



**Note:** There are both IPv4 and IPv6 definitions connecting the full repositories so that definitions for both new and existing cluster definitions are replicated between them. Also be aware that the queue managers QM1 and QM4 cannot communicate directly because they do not share a common network. They could communicate indirectly, for example by using ALIAS queues defined in the queue managers QM2 and QM3. In the configuration shown above you would need to pay attention to the ordering of application messages flowing between QM2 and QM3 because multiple routes exist, if this is relevant you could use BIND\_OPEN to fix the route.

## Abbreviated migration scenarios

This section gives some abbreviated scenarios for when you are thinking of installing clusters on IBM MQ

### Abbreviated scenarios: Effects of CONNAME and LOCLADDR settings

The following table provides an overview of what will occur for the different TCP/IP stacks (IPv4 only, IPv6 only and dual IPv4 and IPv6 stacks) and given the settings for CONNAME and LOCLADDR the expected connection result.

**Note:** Using mapped addresses can require protocol translators in the IP network.

*Table 67. Effects of CONNAME and LOCLADDR settings*

| Stack Type               | CONNAME setting                                    | LOCLADDR setting                                   | Connection result                            |
|--------------------------|----------------------------------------------------|----------------------------------------------------|----------------------------------------------|
| IPv4 only stack          | IPv4 address                                       |                                                    | Channel binds to IPv4 stack                  |
|                          | IPv6 address                                       |                                                    | Channel fails to resolve CONNAME             |
|                          | Host name resolves to both IPv4 and IPv6 addresses |                                                    | Channel binds to IPv4 stack                  |
|                          | IPv4 address                                       | IPv4 address                                       | Channel binds to IPv4 stack                  |
|                          | IPv6 address                                       | IPv4 address                                       | Channel fails to resolve CONNAME             |
|                          | Host name resolves to both IPv4 and IPv6 addresses | IPv4 address                                       | Channel binds to IPv4 stack                  |
|                          | Any address                                        | IPv6 address                                       | Channel fails to resolve LOCLADDR            |
|                          | IPv4 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv4 stack                  |
|                          | IPv6 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Channel fails to resolve CONNAME             |
|                          | Host name resolves to both IPv4 and IPv6 addresses | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv4 stack                  |
| Dual IPv4 and IPv6 stack | IPv4 address                                       |                                                    | Channel binds to IPv4 stack                  |
|                          | IPv6 address                                       |                                                    | Channel binds to IPv6 stack                  |
|                          | Host name resolves to both IPv4 and IPv6 addresses |                                                    | Channel binds to stack determined by IPADDRV |
|                          | IPv4 address                                       | IPv4 address                                       | Channel binds to IPv4 stack                  |
|                          | IPv6 address                                       | IPv4 address                                       | Channel fails to resolve CONNAME             |
|                          | Host name resolves to both IPv4 and IPv6 addresses | IPv4 address                                       | Channel binds to IPv4 stack                  |

Table 67. Effects of CONNAME and LOCLADDR settings (continued)

| Stack Type      | CONNAME setting                                    | LOCLADDR setting                                   | Connection result                                                                                                                                 |
|-----------------|----------------------------------------------------|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
|                 | IPv4 address                                       | IPv6 address                                       | Maps an IPv4 CONNAME to an IPv4 mapped IPv6 address. IPv6 implementations that do not support IPv4 mapped IPv6 addressing fail to resolve CONNAME |
|                 | IPv6 address                                       | IPv6 address                                       | Channel binds to IPv6 stack                                                                                                                       |
|                 | Host name resolves to both IPv4 and IPv6 addresses | IPv6 address                                       | Channel binds to IPv6 stack                                                                                                                       |
|                 | IPv4 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Maps an IPv4 CONNAME to an IPv4 mapped IPv6 address. IPv6 implementations that do not support IPv4 mapped IPv6 addressing fail to resolve CONNAME |
|                 | IPv6 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv6 stack                                                                                                                       |
|                 | Host name resolves to both IPv4 and IPv6 addresses | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv6 stack                                                                                                                       |
|                 |                                                    |                                                    |                                                                                                                                                   |
| IPv6 only stack | IPv4 address                                       |                                                    | Maps an IPv4 CONNAME to an IPv4 mapped IPv6 address. IPv6 implementations that do not support IPv4 mapped IPv6 addressing fail to resolve CONNAME |
|                 | IPv6 address                                       |                                                    | Channel binds to IPv6 stack                                                                                                                       |
|                 | Host name resolves to both IPv4 and IPv6 addresses |                                                    | Channel binds to IPv6 stack                                                                                                                       |
|                 | Any address                                        | IPv4 address                                       | Channel fails to resolve LOCLADDR                                                                                                                 |
|                 | IPv4 address                                       | IPv6 address                                       | Maps an IPv4 CONNAME to an IPv4 mapped IPv6 address. IPv6 implementations that do not support IPv4 mapped IPv6 addressing fail to resolve CONNAME |
|                 | IPv6 address                                       | IPv6 address                                       | Channel binds to IPv6 stack                                                                                                                       |
|                 | Host name resolves to both IPv4 and IPv6 addresses | IPv6 address                                       | Channel binds to IPv6 stack                                                                                                                       |
|                 | IPv4 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Maps an IPv4 CONNAME to an IPv4 mapped IPv6 address. IPv6 implementations that do not support IPv4 mapped IPv6 addressing fail to resolve CONNAME |
|                 | IPv6 address                                       | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv6 stack                                                                                                                       |
|                 | Host name resolves to both IPv4 and IPv6 addresses | Host name resolves to both IPv4 and IPv6 addresses | Channel binds to IPv6 stack                                                                                                                       |

## Abbreviated scenarios: System configurations

Table 69 gives a number of abbreviated scenarios based on the configuration of the installed queue managers and the IP configuration they are running on. The list is not intended to be exhaustive, but to give a number of examples of what to expect based on the configurations shown.

The abbreviations are combined in Table 69 to give the configuration of the systems involved in trying to establish communication. For example:

- v71 + IPv6: Represents a queue manager from an earlier version of the product on a system with a TCP/IP Version 6 stack
- v8 + Dual: Represents a queue manager from the latest version of the product on system with a dual TCP/IP version 4 and Version 6 stack

Table 68. Abbreviations used in system configurations

| Abbreviation | Meaning                                                                                          |
|--------------|--------------------------------------------------------------------------------------------------|
| v71          | queue manager from an earlier version of the product                                             |
| v8           | queue manager from the latest version of the product                                             |
|              |                                                                                                  |
| IPv4         | a system using an IPv4 only stack                                                                |
| IPv6         | a system using an IPv6 only stack                                                                |
| Dual         | a system using both an IPv4 and an IPv6 stack                                                    |
|              |                                                                                                  |
| IPv4DNS      | DNS returns an IPv4 address only for hostname of system holding the responding queue manager     |
| IPv6DNS      | DNS returns an IPv6 address only for hostname of system holding the responding queue manager     |
| DualDNS      | DNS returns an IPv4 and IPv6 address for hostname of system holding the responding queue manager |
|              |                                                                                                  |
| LOCLADDR4    | The LOCLADDR parameter is set to IPv4 addressing                                                 |
| LOCLADDR6    | The LOCLADDR parameter is set to IPv6 addressing                                                 |
|              |                                                                                                  |
| IPADDR4      | IPADDRV is set to IPv4 addressing                                                                |
| IPADDR6      | IPADDRV is set to IPv6 addressing                                                                |

Table 69. System configurations

| Originating queue manager |                            | Responding queue manager |                          |                    | Result                             |
|---------------------------|----------------------------|--------------------------|--------------------------|--------------------|------------------------------------|
| Queue manager and Stack   | LOCLADDR                   | IPADDRV                  | Queue Manager and Stack  | DNS Return         |                                    |
| v71 + IPv6                | Any                        | Not applicable           |                          |                    | IP Error                           |
| v71 + IPv4 or v71 + Dual  | Both LOCLADDR4 & LOCLADDR6 | Not applicable           | v71 + IPv4 or v71 + Dual | IPv4DNS or DualDNS | IPv4 connection can be established |
| v71 + IPv4 or v71 + Dual  | Blank or LOCLADDR4         | Not applicable           | v71 + IPv4 or v71 + Dual | IPv4DNS or DualDNS | IPv4 connection can be established |



Table 69. System configurations (continued)

| Originating queue manager |                                                  | Responding queue manager |                                       |                    | Result                                                                               |
|---------------------------|--------------------------------------------------|--------------------------|---------------------------------------|--------------------|--------------------------------------------------------------------------------------|
| Queue manager and Stack   | LOCLADDR                                         | IPADDRV                  | Queue Manager and Stack               | DNS Return         |                                                                                      |
| v71 + IPv4 or v71 + Dual  | Blank or LOCLADDR4                               | Not applicable           | v71 + Dual                            | IPv6DNS            | Unable to resolve CONNAME                                                            |
| v71 + IPv4 or v71 + Dual  | Blank or LOCLADDR4                               | Not applicable           | v71 + Dual or v8 + Dual<br>v8 + IPv4  | IPv4DNS or DualDNS | IPv4 connection can be established                                                   |
| v71 + IPv4 or v71 + Dual  | LOCLADDR6                                        | Not applicable           |                                       |                    | IP Error                                                                             |
| v71 + IPv4 or v71 + Dual  | Blank or LOCLADDR4 or both LOCLADDR4 & LOCLADDR6 | Not applicable           | v8 + IPv6                             | IPv6DNS            | Unable to resolve CONNAME                                                            |
|                           |                                                  |                          |                                       |                    |                                                                                      |
| v8 + IPv4                 | Blank or LOCLADDR4                               | Not specified            | v71 + IPv4 or v71 + Dual or v8 + IPv4 | IPv4DNS or DualDNS | IPv4 connection can be established                                                   |
| v8 + IPv4                 | LOCADD6                                          | Not specified            |                                       |                    | Unable to resolve LOCLADDR                                                           |
| v8 + IPv4                 | Blank or LOCLADDR4                               | Not specified            | v8 + IPv6                             | IPv6DNS            | Unable to resolve CONNAME                                                            |
| v8 + IPv6                 | Blank or LOCLADDR6                               | Not specified            | v71 + Dual                            | DualDNS            | Attempts to start IPv6 channel and fails as there will be no IPv6 listener available |
| v8 + IPv6                 | Blank or LOCLADDR6                               | Not specified            | v71 + IPv4                            | IPv4DNS            | Attempts to start IPv6 channel and fails as there will be no IPv6 listener available |
| v8 + IPv6 or v8 + Dual    | LOCLADDR6                                        | Blank or IPADDR6         | v8 + IPv6 or v8 + Dual                | IPv6DNS or DualDNS | IPv6 connection can be established                                                   |
| v8 + Dual                 | LOCLADDR6                                        | IPADDR4                  | v8 + Dual                             | IPv4DNS or DualDNS | IPv6 connection can be established where mapped addressing can be used               |
| v8 + Dual                 | Blank or LOCLADDR4                               | IPADDR4                  | v71 + Dual                            | IPv4DNS or DualDNS | IPv4 connection can be established                                                   |
| v8 + Dual                 | Both LOCLADDR4 & LOCLADDR6                       | Blank or IPADDR4         | v71 + Dual                            | IPv4DNS or DualDNS | IPv4 connection can be established                                                   |
| v8 + Dual                 | LOCLADDR4                                        | IPADDR4                  |                                       |                    | Unable to resolve LOCLADDR                                                           |

Table 69. System configurations (continued)


| Originating queue manager |                                                     | Responding queue manager |                           |                       | Result                                |
|---------------------------|-----------------------------------------------------|--------------------------|---------------------------|-----------------------|---------------------------------------|
| Queue manager and Stack   | LOCLADDR                                            | IPADDRV                  | Queue Manager and Stack   | DNS Return            |                                       |
| v8 + Dual                 | LOCLADDR6<br>or both<br>LOCLADDR4<br>&<br>LOCLADDR6 | Blank or<br>IPADDR6      | v8 + IPv6 or<br>v8 + Dual | IPv6DNS or<br>DualDNS | IPv6 connection can be<br>established |

## Overview of migration methods



An overview of the various methods of migrating IBM MQ from one release to another are explained, together with a brief description of the difference between maintenance, migration, and upgrading.

### Maintenance, migration and upgrading

How IBM MQ uses this terminology:

- Migration is the process of updating queue manager data to match a newer level of code. This occurs the first time a queue manager is started with the newer level of code.
- Maintenance is the application of a fix pack, interim fix or program Temporary Fix (PTF). The installation can be restored to its previous level and queue managers or applications continue to work. Migration is not required after applying maintenance. However, you should test applications with the new level of IBM MQ code.
- Upgrading is the process of taking an existing IBM MQ installation and upgrading to a new level of code. Unless the upgrade is applying a fix (and not enabling new function), an upgrade must be followed by migration.
- Once migration has occurred, the queue manager can no longer be started by an earlier code level. Queue manager migration is not reversible  (except on z/OS, with important restrictions).

#### Notes:

1. The example commands in these topics apply to UNIX and Linux, and Windows platforms only.
2.  Side-by-side migration has a different meaning on IBM i. See “Installation methods on IBM i” on page 518 for further information.
3.  Multiple installations are not applicable to IBM i.

### Single stage migration

Single-stage migration is the term used to describe replacing the only installation of IBM MQ on a server, with a later release.

Until IBM WebSphere MQ Version 7.1, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

See “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496 for information on carrying out this method of migration.

## Side-by-side migration

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server.

Queue managers continue running, and remain associated with the older version of IBM MQ, during the installation and verification of the new version of IBM MQ.

When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, migrate them all to the new version, and uninstall the old version of IBM MQ.

The side-by-side migration scenario is less flexible than the multistage migration. The advantage the side-by-side scenario has over the single-stage scenario is that you can install and verify the new IBM MQ installation on the server before switching over to it.

The process is based on the following premise:

- Install additional IBM MQ code alongside existing installation while queue managers are still running.
- Move queue managers one at a time to the new installation.
- Migrate and test applications one at a time.

With the side-by-side approach, you can assign a later version of IBM MQ to be the primary installation, whereas, with the multistage approach, you cannot do so, if you have IBM WebSphere MQ Version 7.0.1 installed, until you uninstall IBM WebSphere MQ Version 7.0.1.

With a later version of IBM MQ set as the primary installation, many applications restart without having to re-configure their environment, and IBM MQ commands work without providing a local search path.

See “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500 for information on carrying out this method of migration.

## Multistage migration

Multistage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server.

After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications.

At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on your staff managing the migration.

See “UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505 for information on carrying out this method of migration.

## Multiple IBM MQ installations

IBM WebSphere MQ Version 7.1 introduced support for multiple IBM MQ installations on Windows, AIX, HP-UX, Linux, and Solaris (but not IBM i ) and gives you the option to install and select between one or more IBM MQ installations.

### z/OS

On z/OS multiple releases can exist, but the information within this topic does not apply to z/OS.

Use STEPLIBs to control which level of IBM MQ is used.

For further information, see "Coexistence" on page 581.

## Overview

You can select between:



- Simplicity of maintaining a single IBM MQ installation.
- Flexibility, by allowing up to a maximum of 128 IBM MQ installations on a system. The first of these 128 installs is reserved for an (optional) IBM WebSphere MQ Version 7.0.1 installation of level 7.0.1.6 or above. The other installations must be IBM WebSphere MQ Version 7.1 or above.

You can install multiple copies of the same code level; this is especially convenient for maintenance purposes.

For example, if you want to upgrade IBM WebSphere MQ Version 7.1.0.0 to IBM WebSphere MQ Version 7.1.0.1, you can install a second copy of IBM WebSphere MQ Version 7.1.0.0, apply the maintenance to bring it to IBM WebSphere MQ Version 7.1.0.1, and then move the queue managers across to the new installation.

You still have the original installation, so it is a simple matter to move the queue managers back if you encounter any problems.

### Notes:

1.   On Linux and Solaris only, you must ensure that each package installed has a unique name.  
You need to use a tool to create a unique set of packages:
  - `$ crtmqpkg PACKAGE_SUFFIX`
  - This takes the IBM MQ installation packages, and repackages them with a new name of your choice. You then install as usual.
2. All installations share a data directory; this is where `mqs.ini` is located for example.
3. All installations share the same namespace for queue managers. This means that you cannot create several queue managers of the same name in different installations.
4. IBM MQ installations are fully relocatable; each installation has a separate installation path. You can choose where you would like to install IBM MQ.
5. IBM MQ resources have installation-scope resource isolation, so operations on one installation do not affect the others.  
This means that the resources created by one installation are isolated from those created by other installations. It enables actions, such as removing an installation of IBM MQ, while queue managers are running under another installation.
6. Queue managers are "associated" with an installation You can move them, but you cannot migrate data back to earlier releases.

## Working with multiple installations

To work with a queue manager, you need to use the commands from its installation. If you select the wrong installation, you see:

```
AMQ5691: Queue manager 'MYQM' is associated with a different installation (Inst1)
```

To work with a queue manager, you have to use the control commands from its associated installation. You have a choice of:

- Using the full path to the control commands, for example:

```
$ MQ_INSTALLATION_PATH\bin\strmqm MYQM
```

or

- Setting the environment variables for an installation with one of:

```
$ MQ_INSTALLATION_PATH/bin/setmqenv 's
$ setmqenv -m MYQM
$ setmqenv -n InstallationName
$ setmqenv -p MQ_INSTALLATION_PATH
```

You might consider using a shell script or batch file to set up the environment for each IBM MQ installation. You can use the **setmqenv** or **crtmqenv** commands to help with this.

- **setmqenv** sets the values of the environment variables, such as **PATH**, **CLASSPATH** and **LD\_LIBRARY\_PATH**, for use with an IBM MQ installation.
- **crtmqenv** creates a list of the environment variables and their values for use with a particular IBM MQ installation. You can then use this list to incorporate into a shell script or batch file.

## Commands

In general, you must use the commands for the correct installation. However, there are some exceptions:

- Commands that work across installations
  - **dspm**
  - **dspm**inst
  - **dspm**ver
  - **setmq**inst
- New control commands for multiple installations
  - **crtmq**env
  - **dspm**inst
  - **setmq**env
  - **setmq**inst
  - **setmq**m

## Primary installation

Explanation of the IBM MQ primary installation.

### Overview

The primary installation is:

- The installation to which system-wide locations refer
- Optional, but convenient

**Note:** Prior to IBM WebSphere MQ Version 7.1, the only installation was the primary installation. However, if you uninstall IBM WebSphere MQ Version 7.0.1, and then install IBM WebSphere MQ Version 7.1, the installation is not by default the primary installation.

### UNIX systems

The primary installation:

- Has symbolic links in `/usr/lib` and `/usr/bin`  
If you have not set a primary installation there are no symbolic links.
- Must be configured manually using the following command:  

```
$ MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

To locate your various installations, you can use the:

- Platform installation tools to query what is installed where on the system
- `dspmqr` command
- `dspmqinst` command
- Following command to list the installations:  

```
cat /etc/opt/mqm/mqinst.ini
```

### Windows

The primary installation is:

- By default the first installation.
- Pointed to by global environment variables.
- Used by some operating system features that require central registration of interface libraries.  
For example, .NET monitor (transactional mode) and COM/ActiveX interface classes.

To locate your various installations, you can use the:

- Platform installation tools to query what is installed where on the system
- `dspmqr` command
- `dspmqinst` command
- Following command to query the registry:  

```
reg.exe query "HKLM\Software\[Wow6432Node\]IBM\WebSphere MQ\Installation" /s
```

### Migration of Windows registry information

IBM WebSphere MQ Version 7.1 onwards uses `mqs.ini` and `qm.ini`.

However, if IBM WebSphere MQ Version 7.0.1 is still installed, IBM WebSphere MQ Version 7.1 uses the registry.

To uninstall IBM WebSphere MQ Version 7.0.1 with IBM WebSphere MQ Version 7.1 already installed, you must:

- End all the queue managers and listeners
- Uninstall IBM WebSphere MQ Version 7.0.1 which migrates the registry information into `mq5.ini` and `qm.ini`.
- IBM WebSphere MQ Version 7.1 now starts using `mq5.ini` instead.

**Note:** If for some reason the uninstall is interrupted, and the automatic migration does not run, a tool is provided to migrate the information out of the registry manually.

## Multiple installations and application programs

When a local application connects to a queue manager, the application needs to load the libraries from the installation associated with the queue manager. Multiple installations introduce some complexity.

### Using the `setmqm` command

When you use `setmqm` to change the installation associated with a queue manager, the libraries that need to be loaded change.

When an application connects to multiple queue managers owned by different installations, multiple sets of libraries need to be loaded.

**Note:** If you link your applications to IBM WebSphere MQ Version 7.1 (or later) libraries, the applications automatically load the appropriate libraries when the application connects to a queue manager.

## Loading IBM MQ libraries in a multi-version environment

How libraries are located depends upon your environment.

If IBM WebSphere MQ Version 7.1 (or later) is installed in the default location, existing applications continue to work as before. Otherwise, you might need to rebuild the application or change your configuration.

The order in which libraries are searched, depends upon the platform you are using:

- Windows
  - The application's directory
  - The current directory
  - The global and your `PATH` variables
- Other platforms
  - `LD_LIBRARY_PATH` (or `LIBPATH/SHLIB_PATH`)
  - An embedded search path (`RPath`)
  - The default library path

Table 70. Options for loading libraries

| Platform     | Option                                                                | Benefits                                                                                                                                  | Drawbacks                                                                                                                                                                                   |
|--------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNIX systems | Set/change the embedded runtime search path (RPath)                   | The path is explicit in the way the application is built                                                                                  | You need to recompile and link<br><br>If you move IBM MQ, you must change RPath                                                                                                             |
| UNIX systems | Set LD_LIBRARY_PATH or equivalent using setmqenv                      | Overrides RPath<br><br>No changes to existing applications<br><br>Easy to change if you move IBM MQ                                       | Depends on environment variables<br><br>Possible impacts on other applications                                                                                                              |
| Windows      | Set PATH using setmqenv                                               | No changes to existing applications<br><br>Easy to change if you move IBM MQ                                                              | Depends on environment variables<br><br>Possible impacts on other applications                                                                                                              |
| All          | Set the primary installation to IBM WebSphere MQ Version 7.1 or later | No changes to existing applications<br><br>Easy to change the primary installation<br><br>Similar behavior to previous versions of IBM MQ | While IBM WebSphere MQ Version 7.0.1 is installed, you cannot make IBM WebSphere MQ Version 7.1 the primary installation<br><br>UNIX systems: Relies on /usr/lib in the default search path |

## Maintenance, upgrade, and migration

Maintenance is a reversible change to the code level of IBM MQ. Maintenance requires no migration. Upgrading is the process of changing the code level of IBM MQ. Migration is the process of updating queue managers, and other objects, such as applications or administrative procedures.

Maintenance is the application of a fix pack, interim fix or PTF. It has one main characteristic. Those fixes, whether applied using a maintenance installation tool, or installed using a manufacturing refresh on top of an installation, are at the same command level as the existing code. No migration is required after applying maintenance. The installation can be restored to its previous level and any changed queue managers or applications will continue to work at the restored code level.

Upgrading and migration are related but distinct. Upgrading is the process of taking an existing IBM MQ installation and upgrading to a new level of code. Unless you are upgrading the fix level of IBM MQ, but not its command level, an upgrade must be followed by migration. Migration is the process of converting queue managers, applications, and other objects to run at a new command level.

An upgrade can take four different forms:

1. Application of a fix pack, interim fix, or a program temporary fix (PTF) using maintenance installation tool. Upgrades applied this way might not be called upgrades, but just fixes. Fixes, applied using a maintenance installation tool, can be rolled back completely as long as no queue manager migration has taken place, and IBM MQ is returned to its previous code level.
2. Installation of new code on top of existing code. You might be able to roll back an upgrade applied in this way; it depends on the platform. Generally speaking, you cannot roll back the installation of new code. To restore the old code level, you must retain the old installation media, and any fixes you applied.




3. Removal of the old level of code, followed by installation of the new level. The installers on very few platforms require you to remove an old installation first. Needless to say, to restore the old code level, you must reinstall it and any fixes.
4. Side by side installation. On z/OS you can install different code levels alongside each other on the same server. In the Job Control Language to start a subsystem, you select the code level to use. On UNIX, Linux, and Windows, you associate a queue manager with an installation, and start the queue manager. In IBM MQ, running multiple queue managers at different command levels on the same server is termed queue manager coexistence. You must not infer from this you can select different installations to run a queue manager at different times. Once a queue manager has been run, it is subject to the rules regarding reverting to earlier or later command levels.

Migration always follows an upgrade that changes the queue manager command level, both automatic and manual changes. Migration is the transformation of queue manager data, applications, and the environment that the queue manager runs in. Migration, and maintenance, and upgrading are described in the following topics.

Upgrades can be backed out, as long as no migration has taken place. The process of removing an upgrade varies by platform and how the upgrade was applied. Upgrades that change the command level of IBM MQ require queue manager migration before applications can reconnect.

  Migration cannot be reversed.

 Migration is a two stage process. Only the first stage, called compatibility mode, is reversible.

## IBM MQ maintenance

Maintenance is the application of a reversible fix. Any changes to queue manager data are compatible with the previous code level.

  Maintenance is the process of applying interim fixes or fix-packs.

 Maintenance is the process of applying program temporary fixes on z/OS.

An important characteristic of applying maintenance is that it must be reversible. Reversibility implies two things:

1. The previous level of code is fully restored.
2. Changes that are made to IBM MQ objects are compatible. Changes are things like the creation or deletion of persistent messages, changes to queue managers, channels, topics, and queues. New and modified objects continue to work correctly with the restored level of code.

The reversibility of a maintenance package limits the extent of functional changes that are included in a maintenance package. No irreversible changes are included in a maintenance package. But, reversibility has limits. A maintenance package might include new programming and administrative interfaces. If you build new or modified applications to use the new interfaces, those applications do not work, if the maintenance package is removed.

Multi-instance queue managers are a good example. Should you remove the Version 7.0.1 fix pack that upgraded Version 7.0, then multi-instance queue manager functions no longer work. However, the queue managers continue to function correctly as single instance queue managers in Version 7.0.

On a smaller scale, a fix pack or interim fix might introduce a new configuration parameter to solve a problem. If you remove the fix pack or interim fix, although the new interface introduced by the change is not available any more, IBM MQ works with any objects that have been changed by the configuration parameter. For example, a new Java system property might introduce a parameter to set a code page for queue manager data conversion. The fix does not change any existing persistent queue manager state

information. It can be removed, and the queue manager continues to work as before, but without the capability introduced in the fix.

#### Related concepts:

 “The version naming scheme for IBM MQ for z/OS” on page 422

On IBM MQ for z/OS, releases have a three-digit Version, Release, and Maintenance (VRM) level code. The code is significant; it identifies the service life of a release. To run a queue manager at a different VRM level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

“The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423


IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

### IBM MQ upgrades and fixes

The term upgrade applies to changing the version V, release R, or modification M of a product. The term fix applies to a change in the F digit.

When you upgrade from one release to another, or apply maintenance refresh packs, fix packs, or interim fixes, the impact of the change depends on the extent of the change in V,R,M,F level. The VRM codes are explained in “The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423.

At each change of V, R, or M, the command level on the queue manager changes, but on a change to F, the command level does not.

 On distributed platforms, after an upgrade has been applied, the only way to *back out* a V.R.M change is by:

- Uninstalling the product code and reinstalling the code, or
- Installing the old level of code alongside the existing code and using the `setmqm` command to associate the queue manager with the other installation.

The general rule, is that if you have carried out an install that causes the command level of the new installation to be updated, and started the queue manager, you cannot *back out* the changes.

### Characteristics of fixes

Application of a fix pack, interim fix, or a program temporary fix (PTF) using a maintenance installation tool should be called a fix.

Fixes, applied using a maintenance installation tool, can be rolled back completely, as long as no queue manager migration has taken place on:

- AIX
- Windows
- z/OS

and IBM MQ is returned to its previous code level.

On all other platforms you must reinstall the product.

### Characteristics of different types of upgrade

An upgrade can take one of three different forms:

1. Installation of new code on top of existing code. You might be able to roll back an upgrade applied in this way; it depends on the platform. Generally speaking, you cannot roll back the installation of new code. To restore the old code level, you must retain the old installation media, and any fixes you applied.

2. Removal of the old level of code, followed by installation of the new level. The installers on very few platforms require you to remove an old installation first. Needless to say, to restore the old code level, you must reinstall it and any fixes.
3. Side by side installation.
  - ▶ **z/OS** On z/OS you can install different code levels alongside each other on the same server. In the Job Control Language to start a subsystem, you select the code level to use.
  - ▶ **Linux** ▶ **UNIX** ▶ **Windows** On UNIX, Linux, and Windows, you associate a queue manager with an installation, and start the queue manager. In IBM MQ, running multiple queue managers at different command levels on the same server is termed queue manager coexistence.

You must not infer from this, that you can select different installations to run a queue manager at different times. Once a queue manager has been run, it is subject to the rules regarding reverting to earlier or later command levels.

**Note:** The term upgrade does not imply that an IBM MQ installation can be directly upgraded from one level to another. On some platforms, an upgrade requires that you remove the previous IBM MQ installation. You can retain any queue managers that you have created.

▶ **z/OS** On z/OS, reversibility of an upgrade has two parts; backout of the installation to the previous code level, and reversion of any queue managers that have been started at the new code level, to work with the previous code level again. See “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452 for more information.

The rules regarding the reversibility of an queue manager to run on a previous code level is dependent on the platform.

▶ **distributed** On IBM i, UNIX, Linux, and Windows, changes in version, release, or modification level are not fully reversible, but changes in fix level are reversible under certain conditions.

An irreversible upgrade implies that you must back up the queue managers, or your system, before upgrading, to be able to restore your queue managers. Taking a backup of a queue manager requires you to stop the queue manager. If you do not take a backup, you are not able to restore IBM MQ to its previous level. Any changes you make on the new level cannot be restored onto the backup system. Changes include the creation or deletion of persistent messages, and changes to queue managers, channels, topics, and queues.

#### Related concepts:

▶ **z/OS** “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

“New function in maintenance level upgrades (On platforms other than z/OS )” on page 451

IBM might introduce new functions between releases in maintenance level upgrades such as fix packs. A maintenance level upgrade including new function increases the maximum command level of an installation.

#### Related reference:

▶ **z/OS** “z/OS: OPMODE” on page 677

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. To access V8.0 capabilities, change the value of **OPMODE** to **OPMODE=(NEWFUNC,800)**. To restrict the use of new capabilities, and retain the ability to revert the queue manager to its earlier level, leave **OPMODE** at its default setting, **OPMODE=(COMPAT,800)**.

## Windows - upgrading an IBM MQ installation:

To upgrade an IBM MQ installation on Windows, from one version, release, and modification level to a later one, carry out the tasks detailed in one of the following subtopics.

Before you begin, ensure that you have backed up your data.

**Important:** If you want to apply maintenance instead, see the Windows topics listed in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.

### Related tasks:

“Migrating a queue manager to the latest release” on page 512

The procedures for migrating a queue manager to the latest release are detailed in the following topics.

### *Upgrading an IBM MQ server installation using the Launchpad:*

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using the Launchpad.

### Before you begin

Ensure that you have:

1. Stopped all your IBM MQ applications
2. Shut down your listeners
3. Stopped all your queue managers
4. Backed up your data

**Important:** If you want to apply maintenance instead, see the Windows topics listed in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.

### Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate `setup.exe` in the base directory of the IBM MQ installation image.
  - From a DVD, this location might be:  
`E:\setup.exe`
  - From a network location, this location might be:  
`m:\instmq\setup.exe`
  - From a local file system directory, this location might be:  
`C:\instmq\setup.exe`
3. Double-click the **Setup** icon to start the installation process. It is possible to run either by:
  - Running `setup.exe` from the command prompt. Or
  - Double-clicking `setup.exe` from Windows Explorer.

If you are installing on a Windows system with UAC enabled, accept the Windows prompt to allow the launchpad to run as elevated. During installation, you might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the installation to continue.

The IBM MQ Installation Launchpad window is displayed.

4. Continue to follow the Launchpad instructions as shown on screen.

5. Select **Installing a new instance**, if you see a panel asking you to choose between installing a new instance, or maintaining or upgrading an existing instance, when you click the **Launch IBM MQ Installer** button. You use the other option when adding or removing features from an already installed IBM MQ.
6. On the next panel, choose between **Install leaving the existing installation(s) untouched** or **Upgrade an existing named installation already on the machine**, and click **Next**.  
**Attention:** If you do not see this screen, it means that there was no IBM MQ server installation on the machine that could be upgraded by this installer.
7. Follow the installer prompts to upgrade your IBM MQ server installation.

**Related tasks:**

“Upgrading an IBM MQ server installation using msixec” on page 446

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using msixec.

“Upgrading an IBM MQ client installation using the Launchpad”

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using the Launchpad.

“Upgrading an IBM MQ client installation using msixec” on page 447

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using msixec.

*Upgrading an IBM MQ client installation using the Launchpad:* 

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using the Launchpad.

**Before you begin**

Ensure that you have:

1. Stopped all your IBM MQ applications
2. Shut down your listeners
3. Stopped all your queue managers
4. Backed up your data

**Procedure**

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate setup.exe in the base directory of the IBM MQ installation image.
  - From a DVD, this location might be:  
`E:\setup.exe`
  - From a network location, this location might be:  
`m:\instmq\setup.exe`
  - From a local file system directory, this location might be:  
`C:\instmq\setup.exe`
3. Double-click the **Setup** icon to start the installation process. It is possible to run either by:
  - Running setup.exe from the command prompt. Or
  - Double-clicking setup.exe from Windows Explorer.

If you are installing on a Windows system with UAC enabled, accept the Windows prompt to allow the launchpad to run as elevated. During installation, you might also see Open File - Security Warning dialog boxes that list International Business Machines Limited as the publisher. Click **Run** to allow the installation to continue.

The IBM MQ Installation Launchpad window is displayed.

4. Continue to follow the Launchpad instructions as shown on screen.
5. Select **Installing a new instance**, if you see a panel asking you to choose between installing a new instance, or maintaining or upgrading an existing instance, when you click the **Launch IBM MQ Installer** button. You use the other option when adding or removing features from an already installed IBM MQ.
6. On the next panel, choose between **Install leaving the existing installation(s) untouched** or **Upgrade an existing named installation already on the machine**, and click **Next**.  
**Attention:** If you do not see this screen, it means that there was no IBM MQ client installation on the machine that could be upgraded by this installer.
7. Follow the installer prompts to upgrade your IBM MQ client installation.

#### Related tasks:

“Upgrading an IBM MQ server installation using the Launchpad” on page 444


How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using the Launchpad.

“Upgrading an IBM MQ client installation using the Launchpad” on page 445

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using the Launchpad.

“Upgrading an IBM MQ client installation using msixec” on page 447

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using msixec.

Upgrading an IBM MQ server installation using msixec: 

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using msixec.

#### Before you begin

Ensure that you have:

1. Stopped all your IBM MQ applications
2. Shut down your listeners
3. Stopped all your queue managers
4. Backed up your data

**Important:** if you want to apply maintenance instead, see the Windows topics listed in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.

#### Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate MSI file in the MSI directory of the IBM MQ installation image.
  - From a DVD, this location might be:  
`E:\MSI\IBM MQ.msi`
  - From a network location, this location might be:  
`m:\instmq\MSI\IBM MQ.msi`
  - From a local file system directory, this location might be:  
`C:\instmq\MSI\IBM MQ.msi`
3. Optional: If you are upgrading the only IBM MQ server installation, where the installation has the default value `Installation1` issue the following command:

```
msiexec /i "<InstallationImage>\MSI\IBM MQ.msi" /q AGREETOLICENSE=YES
INSTALLATIONNAME="Installation1"
```

4. Optional: If you are upgrading an installation on a machine that already has one or more IBM MQ server installations of the level you are upgrading to, you must provide additional parameters to select a free MSI instance ID. See [Choosing MSI Instance IDs for multiple server installations](#) for more information.

In this case, the command might look something like:

```
msiexec /i "<Installation Image>\MSI\IBM MQ.msi" /q AGREETOLICENSE=YES
INSTALLATIONNAME="Installation2" NEWINSTANCE=1
TRANSFORMS=":InstanceId2.mst;1033.mst"
```

#### Related tasks:

[“Upgrading an IBM MQ server installation using the Launchpad” on page 444](#)

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using the Launchpad.

[“Upgrading an IBM MQ client installation using the Launchpad” on page 445](#)

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using the Launchpad.

[“Upgrading an IBM MQ client installation using msiexec”](#)

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using msiexec.

*Upgrading an IBM MQ client installation using msiexec:* 

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using msiexec.

#### Before you begin

Ensure that you have:

1. Stopped all your IBM MQ applications
2. Shut down your listeners
3. Stopped all your queue managers
4. Backed up your data

#### Procedure

1. Access the IBM MQ installation image. The location might be the mount point of the DVD, a network location, or a local file system directory.
2. Locate MSI file in the MSI directory of the IBM MQ installation image.
  - From a DVD, this location might be:  
`E:\MSI\IBM MQ.msi`
  - From a network location, this location might be:  
`m:\instmq\MSI\IBM MQ.msi`
  - From a local file system directory, this location might be:  
`C:\instmq\MSI\IBM MQ.msi`
3. Optional: If you are upgrading the only IBM MQ client installation, where the installation has the default value `Installation1` issue the following command:

```
msiexec /i "<InstallationImage>\MSI\IBM MQ.msi" /q AGREETOLICENSE=YES
INSTALLATIONNAME="Installation1"
```



4. Optional: If you are upgrading an installation on a machine that already has one or more IBM MQ client installations of the level you are upgrading to, you must provide additional parameters to select a free MSI instance ID. See [Choosing MSI Instance IDs for multiple client installations](#) for more information.

In this case, the command might look something like:

```
msiexec /i "<Installation Image>\MSI\IBM MQ.msi" /q AGREETOLICENSE=YES
INSTALLATIONNAME="Installation2" NEWINSTANCE=1
TRANSFORMS=":InstanceId2.mst;1033.mst"
```

#### **Related tasks:**

[“Upgrading an IBM MQ server installation using the Launchpad” on page 444](#)

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using the Launchpad.

[“Upgrading an IBM MQ client installation using the Launchpad” on page 445](#)

How you upgrade an IBM MQ client installation on Windows to a newer version, release, or modification, using the Launchpad.

[“Upgrading an IBM MQ server installation using msiexec” on page 446](#)

How you upgrade an IBM MQ server installation on Windows to a newer version, release, or modification, using msiexec.

## **IBM MQ migration**

Migration is the conversion of programs and data to work with a new code level of IBM MQ. Some types of migration are required, and some are optional. Queue manager migration is never required after applying a maintenance level update, that does not change the command level. Some types of migration are automatic, and some are manual. Queue manager migration is typically automatic and required after releases and manual and optional after a maintenance level upgrade that introduces a new function. Application migration is typically manual and optional.

Whenever you upgrade IBM MQ to a new release that changes its command level, migration is performed by the queue manager. Whenever you upgrade IBM MQ to a new maintenance or fix level, which introduces a new function using a new command level, you can migrate the queue manager to use the new command level and thereby the new function.

You must read [“Changes that affect migration” on page 580](#) before upgrading your IBM MQ installation or migrating your queue managers, to identify what migration tasks you must plan for.

Using the model in [Figure 58 on page 450](#) you can distinguish different migration questions, which are discussed in the following topics:

### **Operating environment migration**

Upgrading the operating environment, or components in the environment such as installing a new level of JRE; see [“IBM MQ operating environment migration” on page 455](#)

### **Queue manager migration**

Migrating a queue manager following an upgrade of the IBM MQ installation to a new command level; see [“Queue manager migration” on page 472](#).

### **IBM MQ MQI client migration**

Migrating a client configuration following installation of a new version or release of the IBM MQ MQI client ; see [“IBM MQ MQI client migration” on page 474](#).

### **Application migration**

Relinking, recompiling, or recoding an IBM MQ server or client application; see [“Application migration and interoperation” on page 477](#). Application migration also includes migrating any API or channel exits

In addition, you must consider the impact of migrating one queue manager, or IBM MQ MQI client, on other clients or queue managers:



## Compatibility, coexistence, and interoperability

See “Coexistence, compatibility, and interoperability” on page 456 for information about the compatibility of IBM MQ applications connected to queue managers and IBM MQ MQI client clients on different command levels. The section also explains the concept of queue manager coexistence, and the interoperability of IBM MQ JMS applications with WebSphere Application Server.

## Queue manager clusters

Can a queue manager cluster contain queue managers at different command levels? See “Queue manager cluster migration” on page 478 to answer this question, and how to migrate a cluster of queue managers.

### z/OS Queue-sharing groups

Queue sharing groups involve multiple queue managers running on z/OS. How do you migrate queue managers that are part of a queue-sharing group to a new command level; see “Queue-sharing group migration” on page 480.

## High-availability clusters

How do you migrate queue managers that are part of a high-availability cluster to a new command level, and maintain continuous and reliable service? See “Migrating a queue manager in a high-availability configuration” on page 481, which covers both migration of multi-instance queue managers, and the migration of queue managers operating in high-availability clusters.

## IBM MQ migration concepts

Figure 58 on page 450 shows two runtime operating system environments. One environment is called *Server*, and contains an IBM MQ server and server application. The other is called *Client*, and contains an IBM MQ MQI client application. The server environment has one or more queue managers represented by **QM** using the installation of IBM MQ installed on the server.

The queue manager labeled QM-n? coexists on the same server as QM, but runs at a different release level. Multiple releases of IBM MQ installed in the same operating environment are called coexistent<sup>2</sup>. The IBM MQ installations for different release levels are not shown. The question-mark in the queue manager name indicates this capability might not be present in your environment.

z/OS Only z/OS supports multiple queue managers coexisting at different release levels in the same operating environment.

Queue manager coexistence is important for migration in two respects:

1. It can be used to reduce the risk involved in migrating to a new command level, and reduce the downtime during the migration process.
2. You must consider any configuration implications of running some applications or clusters on the same server with queue managers at different command levels.

For details, see “Queue manager coexistence in Version 8.0” on page 457

The queue manager, QM\*, represents queue managers of various levels installed on other servers.

---

2. It is not necessary, but it is usual, for coexistent installations to be at different release levels.

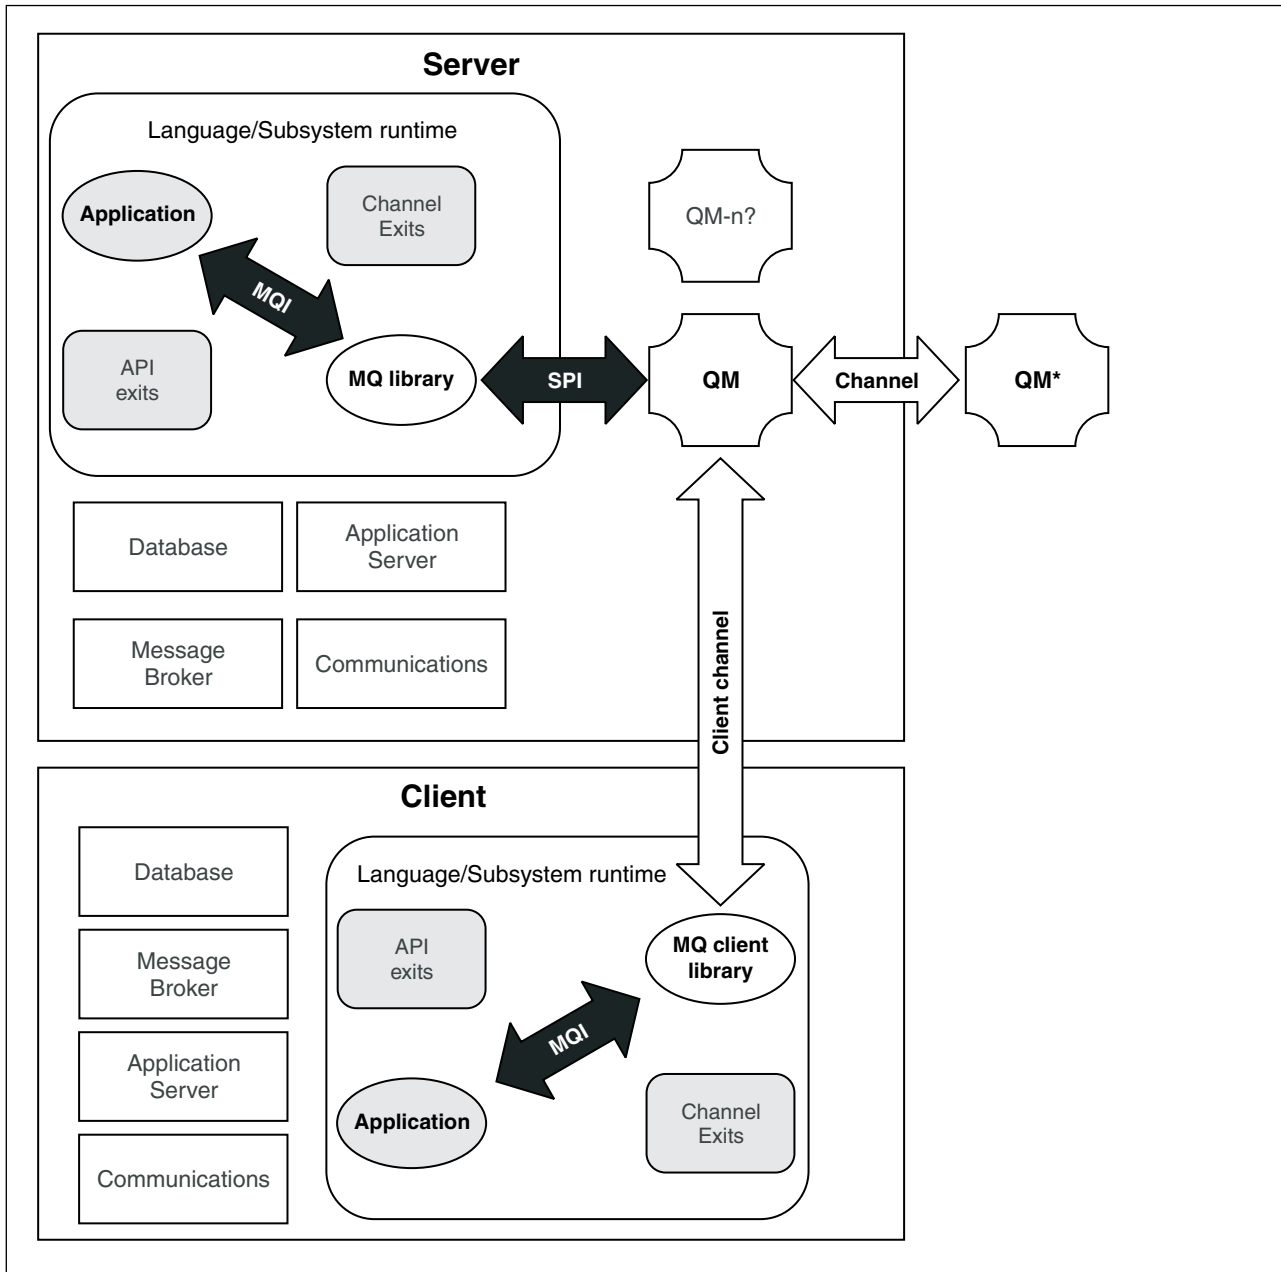


Figure 58. IBM MQ application migration model

## New function in maintenance level upgrades (On platforms other than z/OS )

IBM might introduce new functions between releases in maintenance level upgrades such as fix packs. A maintenance level upgrade including new function increases the maximum command level of an installation.

When you apply the maintenance level upgrade, the installation supports the new command level. A queue manager associated with the installation is not able to use the new function until you have run **strmqm** to set the queue manager to the new command level.

A fix pack introducing a new function also includes regular fixes as a normal fix pack. If you do not want to use the function, but apply the fixes, do not enable the new function for any queue manager. The fix pack then is a fully reversible maintenance level upgrade. If you want to use a new function with a queue manager, you must change the command level of the queue manager. If one queue manager associated with an installation enables the new function, the fix pack introducing the new function and all previous fix packs cannot be removed anymore.

By applying a maintenance level upgrade that introduces a new function the IBM MQ installation supports an additional command level. An installation can therefore support multiple command levels and multiple queue managers associated with it can run on different command levels.

To enable a new function for a queue manager, you must set the command level of queue manager to the command level supporting the new function. Because command levels are cumulative, the command level also supports all other functions introduced by previous maintenance level upgrades. By enabling a new function from a maintenance level upgrade you therefore enable all new functions from previous maintenance level upgrades, too.

### Example

You have installed two fix packs including new functions:

1. Fix pack 7.1.0.2, introducing the command level 711.
2. Fix pack 7.1.0.15, introducing the command level 712.

You might have also installed other fix packs, but these did not introduce new function.

You want to use the new function from fix pack 7.1.0.15 on one of your queue managers. You enable command level 712 for that queue manager using the **strmqm** command. This enables all of the function introduced in both fix packs. It is not possible to use the function introduced by command level 712 without also enabling the function introduced by command level 711.

Once you have enabled command level 712 for a queue manager, the minimum version of product code to start the queue manager is 7.1.0.15. If you uninstall fix pack 7.1.0.15 from the installation with which the queue manager is associated, you will not be able to start the queue manager. In this situation, the **strmqm** command fails and issues error message AMQ7204 indicating that the queue manager has already been started by a newer release. If you have multiple installations of IBM MQ, you can start the queue manager using another installation which satisfies this minimum version requirement.

## Related tasks:

“Migrating queue managers to new-function fix packs” on page 655

This scenario illustrates running different levels of queue manager from a single installation using new-function fix packs. It contrasts migrating a queue manager to new command levels in new-function fix packs, to migrating a queue manager to a new command level in a new release. The scenario explains the relationship between new-function fix packs and maintenance fix packs.

## Upgrade, migration, and maintenance of IBM MQ on z/OS

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

Applying PTFs does not change the version, release, or maintenance level of the code. No queue manager migration is required after applying maintenance. PTFs are grouped into Recommended Service Updates (RSUs) that have been tested together in a Consolidated Service Test (CST); see Consolidated Service Test and the RSU.

When you install a new VRM level of IBM MQ on z/OS using SMP/E, it creates a set of IBM MQ libraries. The libraries for different VRM levels of IBM MQ can coexist on the same instance of z/OS. You can then run different queue managers against different release levels of IBM MQ on the same z/OS instance.

If you start a queue manager running on a later release level, then migration of the queue manager to that release level is required. Even if the difference is only in the modification level, some migration might be required. The migration tasks you must perform to migrate from one version to another are documented in “z/OS: Migration planning to the latest release” on page 486; see also “Changes that affect migration” on page 580.

From Version 7.0.1, after you have fully migrated a queue manager to a new version or release, reverse migration is not possible. Before Version 7.0.1, full migration took place when you started a queue manager for the first time at the new release level. For Version 7.0.1 and later versions, you have control over when migration takes place, using a new **CSQ6SYSP** parameter, **OPMODE**; see “z/OS: OPMODE” on page 677. If your queue manager is on Version 7.0 or earlier, you can revert to an earlier release. You might have to contact your IBM support center for a backward migration PTF.

Using **OPMODE**, you can migrate all your existing applications to the new release level, and still be able to revert to the previous release level. Once you start changing applications, or adding applications that use new function, you cannot revert to the previous level of IBM MQ. **OPMODE** applies to migration from Version 6.0 to Version 7.0.1, onwards.

**OPMODE** gives you the option of enforcing a two-stage migration process:

1. Regression test your existing applications.
2. Develop new applications, and change existing applications, to use the new function in the release.

The strategy for upgrading queue managers at Version 6.0 or later is as follows:

1. Apply the coexistence and backward migration PTFs to all the queue managers you are going to upgrade. After applying the PTFs, you can run queue managers of different levels in the same queue sharing groups. You can also reverse the migration of a queue manager back to your current level.
2. Upgrade the first queue manager.
3. Check all your existing applications run correctly on this queue manager.
4. Bring all the queue managers in a queue sharing group up to the new level, and check that existing applications continue to work correctly.

5. Change the setting of **OPMODE** so that applications can use new function on all the queue managers in the queue sharing group.

**Note:** Step 5 is the point of no return. You can no longer run that queue manager at the previous level of IBM MQ.

6. To enable new IBM MQ v7.1, or later, function, restart all queue managers within the queue sharing group.

The coexistence and backward migration PTFs have two distinct purposes:<sup>3</sup>

1. To allow queue managers at the earlier release level to coexist with ones at the later release level. In particular for queue managers to coexist in the same queue sharing group.
2. To handle queue manager data and logs formatted using the data definitions of the later release.

On z/OS, maintenance is supplied as Program Temporary Fixes, PTFs, which are applied and removed using SMP/E. PTFs are specific to a particular set of libraries corresponding to specific release level. Apart from any exceptions documented with the PTFs, PTFs do not change the correct operation of IBM MQ. Nonetheless, you must check that the fixes have not changed the operation of critical programs unexpectedly.

PTFs that apply to a category of software fixes might be grouped together and identified using a fix category. For further information, see IBM Fix category values and descriptions.

## Characteristics of different types of upgrade on z/OS

When you upgrade from one release to another, or apply maintenance in the form of PTFs, on z/OS, the impact of the change depends on the extent of the change in VRM level. The VRM codes are explained in “The version naming scheme for IBM MQ for z/OS” on page 422.


PTF upgrades do not require migration, and are reversible. From Version 7.0.1, all upgrades from Version 6.0 or later are reversible if the **OPMODE** has not been set to NEWFUNC.

### Related concepts:

“The version naming scheme for IBM MQ for z/OS” on page 422

On IBM MQ for z/OS, releases have a three-digit Version, Release, and Maintenance (VRM) level code. The code is significant; it identifies the service life of a release. To run a queue manager at a different VRM level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

## Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

On different platforms, you employ different mechanisms to install and maintain software releases. Installing a release at a new maintenance level, and applying maintenance level upgrades to upgrade an earlier release to the same maintenance level, have different results.

---

3. Coexistence and backward migration changes might be shipped as a single or multiple fixes.

If you start a queue manager running on a later release level, then migration of the queue manager to that release level is required. The migration tasks you must perform to migrate from one release to another are documented in “Migrating a queue manager to the latest release” on page 512 ; see also “Changes that affect migration” on page 580.

When you upgrade the maintenance or fix level of IBM MQ by applying a regular maintenance level upgrade, you can reverse the upgrade by removing the fix. When you upgrade the maintenance or fix level of IBM MQ by applying a maintenance level upgrade containing a new function, you can reverse that upgrade and all previously reversible upgrades until a queue manager associated with the installation enables the new function.

You cannot easily revert to a previous level of IBM MQ after installation. If you install a copy of IBM MQ obtained from Passport Advantage or from physical media, the installer uninstalls IBM MQ, if it is present. It then installs the new level of IBM MQ. To revert to the previous level of IBM MQ, you must keep the earlier installation image and any fixes you applied. Then you must uninstall the new level, reinstall the previous release level, and reapply the required fixes. If you have started any queue managers at the later level, they will not work with the restored level of IBM MQ<sup>4</sup> . To restore IBM MQ to its previous level, after starting any queue managers, you must first back up the queue managers. You can then restore the backup queue managers after restoring the previous level of IBM MQ.

Maintenance levels and fix levels are both supplied from the service site, Fix Central. Fix central has a function to tell you what upgrades you can apply to the current level of your system. If you back out a maintenance level upgrade, it returns IBM MQ code to the same level of code as before applying the maintenance level upgrade.

### **Characteristics of different types of upgrade**

When you upgrade from one release to another, or apply maintenance refresh packs, fix packs, or interim fixes, the impact of the change depends on the extent of the change in VRMF level. The VRM codes are explained in “The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423.

Table 71 on page 455 describes characteristics of different upgrade paths for IBM MQ. In particular, notice that migration is required only if the version or release number changes. Other types of upgrade do not require migration, and are reversible if the upgrade is applied using a maintenance procedure.

---


4. Unless you installed a later maintenance level upgrade, not a new release or version: then you could revert to an earlier maintenance level by reinstalling the earlier maintenance level upgrade. Queue manager data is compatible between maintenance levels.

Table 71. Types of upgrade (Platforms other than z/OS )

| Upgrade type                                     | Examples                        | Suggested testing                     | Type of installation                                  | Manufacturing refresh | New features | IBM MQ Migration |
|--------------------------------------------------|---------------------------------|---------------------------------------|-------------------------------------------------------|-----------------------|--------------|------------------|
| Version                                          | 7.0                             | Full testing of all applications      | Full or upgrade installation                          |                       | Yes          | Yes              |
| Modification                                     | 7.0.1                           | Regression test critical applications | Full, or upgrade installation, or maintenance package | Yes                   | Minor        |                  |
| Enabled new-function fix pack                    | No example yet                  | Regression test critical applications | Maintenance package & queue manager migration         |                       |              |                  |
| No-function or not enabled new-function fix pack | 7.0.1.3                         | Brief test of critical applications   | Maintenance package                                   | No                    | No           | No               |
| Interim fix                                      | 7.0.0.1-WS-MQ-Windows-LAIZ50784 | Test affected applications            | Manual                                                |                       |              |                  |

#### Related concepts:

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

“The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423

IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

#### Related information:

Backing up and restoring a queue manager

### IBM MQ operating environment migration

You might perform some migration tasks for IBM MQ as a result of upgrading the operating environment.

To find out what operating environment upgrades you must make before upgrading IBM MQ, compare the requirements for different IBM MQ releases. For example, compare List of Supported Software for IBM MQ V7.0 with System Requirements for IBM MQ V8.0.

Note that the System Requirements page for IBM MQ Version 8.0 uses the Software Product Compatibility Reports (SPCR) tool.

By selecting the appropriate link on the web page, the SPCR tool enables you to go directly to the:

- Supported operating systems
- Prerequisites
- System requirements, and
- Optional supported software

for the specific operating system, or systems, that your enterprise uses.

If an operating environment change directly affects the migration to a new version of IBM MQ, it is listed in “Changes that affect migration” on page 580.

The change might affect IBM MQ migration indirectly. For example, the runtime linkage conventions for applications, or the way memory is allocated, might change.

## **Coexistence, compatibility, and interoperability**

The definitions of the IBM MQ terms coexistence, compatibility, and interoperability.

### **Coexistence**

Is being able to install and run two or more versions of the same program on the same server.

For IBM MQ, it normally means installing and running multiple versions of IBM MQ on a server.

### **Compatibility**

Is the ability to run applications from one level of queue manager with an earlier, or previous level, of the queue manager.

If you are using a message channel agent (MCA) channel, any version and release of an IBM MQ queue manager can connect, using an MCA channel, to any version and release of another IBM MQ queue manager.

The MCA channel is automatically configured to the latest version of protocol that is supported by both ends of the channel.

Compatibility is also the ability to run client applications with different versions of the IBM MQ MQI client, and different levels of the queue manager.

### **Interoperability**

Is mainly the ability to exchange messages between different versions of IBM MQ. It can also mean the interoperability between others things, such as publish/subscribe brokers, or between components such as the IBM MQ classes for JMS and WebSphere Application Server.


Maintaining the compatibility, coexistence, and interoperability of IBM MQ is important in order to preserve the investment you make in applications and administrative procedures.

Three areas to which this objective does not apply to as rigidly, are:

- GUI interfaces, such as IBM MQ Explorer.
- Information for service, such as FFST files and traces.
- Error messages. The text in an error message might change, to make the wording clearer or more accurate.



## Queue manager coexistence in Version 8.0

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

### Single installation queue manager coexistence on all platforms

Single installation queue manager coexistence is useful in development and production environments. In development environments, you can set up different queue manager configurations to support different development activities. You can also work with multiple queue manager configurations on a single server, connected by channels, as if deployed on a network.

In production environments configuring multiple queue manager on a single server is less common. It has no performance or functional advantage over a single queue manager configuration. Sometimes, you must deploy multiple queue managers on server. It might be essential to meet the requirements of a particular software stack, governance, administration, or as a consequence of the consolidation of servers.

### Multi-installation queue manager coexistence

Multi-installation<sup>5</sup> queue manager coexistence became available in Version 7.1 on UNIX, Linux, and Windows. Multi-installation queue manager coexistence has always been supported on z/OS.

With multi-installation queue manager coexistence on the same server you can run queue managers at different commands levels on the same server. You can also run multiple queue managers at the same command level, but associate them with different installations.

Multi-installation adds more flexibility to the coexistence of queue managers using a single installation. Any of the reasons behind running multiple queue managers, such as supporting different software stacks, might require different versions of IBM MQ.

The biggest benefit of multi-installation identified by early users, is in upgrading from one version of IBM MQ to another. Multi-installation makes upgrading less risky, less costly, and is more flexible in meeting the migration needs of applications running on a server.

The key to migration flexibility is being able to install a new version alongside an existing installation; see Figure 59, which is extracted from “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500.




Figure 59. Side-by-side installation - step 2

When the installation is complete, and verified, migrate queue managers and applications to the new installation; see Figure 60 on page 458. When migration is complete, uninstall the old installation.

5. Do not confuse multi-installation queue manager coexistence with multi-instance queue managers. They are completely different, though they sound similar in English.




Figure 60. Side-by-side installation - step 4

Think of multi-installation as being the basis for a range of migration strategies. At one end is “Single-stage”, in which you only have one installation on a server at a time. At the other end is multi-stage migration, in which you continue to run multiple installations at the same time. In the middle is side-by-side migration. Each of the three strategies is explained in these three tasks:

1. “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496
2. “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500
3. “UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Another similar use of multi-installation is to support the migration of queue managers to a new fix level; see Figure 61 on page 459. You maintain two installations, one of which has the latest fix pack applied, and the other has the previous maintenance levels. When you have moved all queue managers to the latest fix pack level, you can replace the previous fix pack with the next fix pack to be released. The configuration allows you to stage the migrating applications and queue managers to the latest fix pack level. You can switch the primary installation designation to the latest fix pack level.

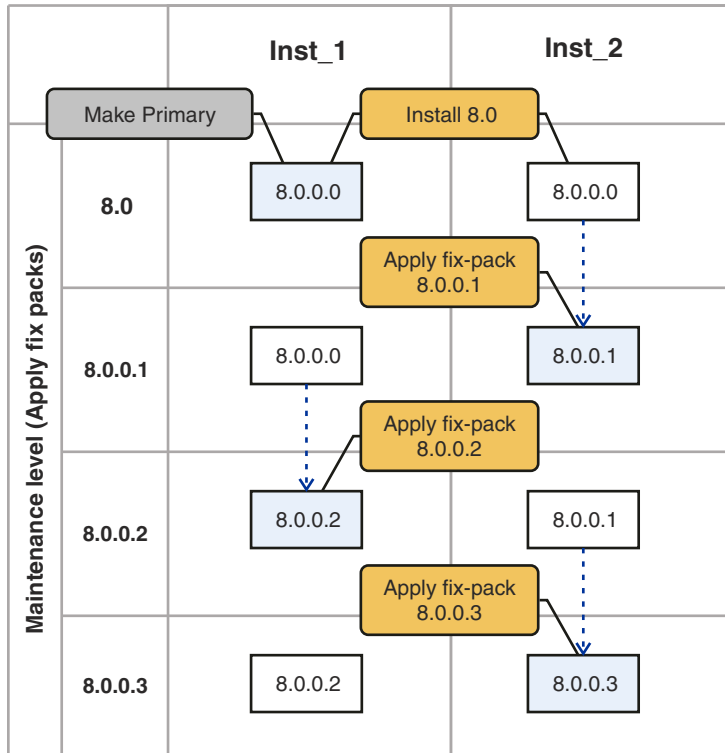


Figure 61. Rolling fix packs

#### Related concepts:

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460  
 You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

**z/OS** “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452  
 You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

#### Related tasks:

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533  
 No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.  
 “UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542  
 Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

“UNIX, Linux, and Windows: Staging maintenance fixes” on page 669

Use multiple installations of IBM MQ on the same server to control the release of maintenance fixes.

“Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 540

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

## Multi-installation queue manager coexistence on UNIX, Linux, and Windows

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

Figure 62 shows two IBM MQ installations, two queue managers, and three applications. Applications 2 and 3 are connected to QM2, and application 1 is connected to QM1. Applications 1 and 3 load IBM MQ libraries from the Inst\_1 installation, and application 2 loads libraries from the Version 7.0.1 installation.




Figure 62. Coexistence of two queue managers using Version 7.0.1 and a later version installations

When you upgrade from Version 7.0.1 to a later version, you can choose to run Version 7.0.1 alongside the later version. The installation, illustrated in Figure 62, is called a multi-version installation. You can also install multiple copies of Version 7.1 alongside each other. That would be called multi-installation. Multi-installation is the more general term.

Version 7.0.1 did not support multi-installation on distributed platforms. Before Version 7.1 becoming available, fix pack 7.0.1.6 was shipped with some fixes to make Version 7.0.1 compatible with a later version on the same server. With 7.0.1.6 installed, you can run one copy of Version 7.0.1 alongside multiple copies of the later version. You do not have to apply the fix pack to upgrade Version 7.0.1 to Version 7.1 “in-place” ; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496.

A multi-version installation that includes Version 7.0.1, does not behave the same way as a multi-installation that does not. The differences primarily affect how you might choose to configure how applications load IBM MQ libraries, and run IBM MQ commands. Because of these differences, think of the multi-version support provided in 7.0.1.6, as a migration aid to moving to a later version multi-installation environment. The topics that explain the restrictions in Version 7.0.1 multi-version are listed in related links.

If you run multiple installations of IBM MQ on a server you must consider three questions:

1. Which installation is a queue manager associated with; see “Queue manager association”?
2. Which installation does an application load; see “Loading IBM MQ libraries”?
3. Which installation is an IBM MQ command run from; see “Command association” on page 463?

### **Queue manager association**

Before Version 7.1, queue managers on UNIX, Linux, or Windows were associated with the only installation on the server. With Version 7.1, or later, installed on the same server as Version 7.0.1, you can change the association of a queue manager to a later version by running **setmqm** ; see **setmqm**. You cannot change the association of a queue manager running a release of IBM MQ earlier than Version 7.0.1 because you cannot install a later version of the product on a server with an installation of IBM MQ earlier than Version 7.0.1.

A queue manager is permanently associated with an installation, until you choose to change the association with the **setmqm** command. You cannot associate a queue manager with an installation at a lower command level than the current command level of the queue manager.

In Figure 62 on page 460, QM1 is associated with Inst\_1. The association is made by running **setmqm -m QM1 -n Inst\_1**. When QM1 is first started, after running **setmqm**, if QM1 was running Version 7.0.1, it is migrated to a later version. QM2 is associated with Version 7.0.1 because the association has not been changed.

### **Loading IBM MQ libraries**

The application connections to the queue managers are established by calling MQCONN or MQCONNX in the normal way.

Which IBM MQ library an application loads depends on the configuration of the operating system loader and on the IBM MQ installation the queue manager is associated with.

In Figure 62 on page 460, the operating system loads the IBM MQ library from the Inst\_1 installation for applications 1 and 3. It loads the IBM WebSphere MQ Version 7.0.1 library for application 2. The operating system has loaded the wrong library for application 3. Application 3 requires the IBM WebSphere MQ Version 7.0.1 libraries.

Figure 63 on page 462 shows what happens to application 3. Application 3 is connecting to QM2, and QM2 is associated with the IBM WebSphere MQ Version 7.0.1 installation. IBM MQ detects that the operating system has loaded the wrong library to process calls from application 3 to QM2. IBM MQ loads the correct library from the IBM WebSphere MQ Version 7.0.1 installation. It transfers the MQCONN or MQCONNX

call to the IBM WebSphere MQ Version 7.0.1 library. Subsequent MQI calls that use the connection handle returned by MQCONN or MQCONNX, call entry points in the IBM WebSphere MQ Version 7.0.1 library.

Because IBM WebSphere MQ Version 7.0.1 libraries cannot load IBM MQ libraries from other installations, there is no corresponding application in Figure 63 that loads a IBM WebSphere MQ Version 7.0.1 library and connects to a queue manager running Version 7.1. If you attempt a connection to QM1 with application 2, IBM MQ returns an error; see 2059 (080B) (RC2059): MQRC\_Q\_MGR\_NOT\_AVAILABLE.

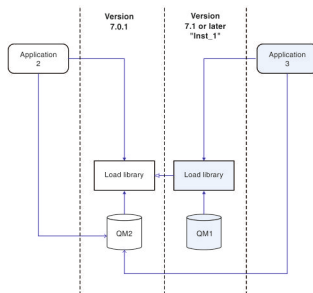


Figure 63. Loading calls in a different library

A Version 7.1, or later, IBM MQ library includes a routing capability that is based on the installation a queue manager is associated with. Earlier IBM MQ libraries do not have a routing capability. The operating system can load a library from any Version 7.1 installation, or later, and IBM MQ transfers MQI calls to the correct library.

The new loading capability of IBM MQ libraries in Version 7.1, or later, does not relax the restriction that an application compiled and linked at a later release level must not directly load an IBM MQ library at an earlier release level. In practice the restriction is of less significance than in earlier releases, because as long as the operating system loads a library at the same or later level than the library the application was compiled and linked with, IBM MQ can call any other level of IBM MQ on the same server from Version 7.0.1 upwards.

For example, suppose you recompile and link an application that is to connect to a Version 7.0.1 queue manager using the libraries shipped with Version 7.1. At run time the operating system must load the Version 7.1 libraries for the application, even though the application connects to a Version 7.0.1 queue manager. IBM WebSphere MQ Version 7.1 detects the inconsistency and loads the Version 7.0.1 library for the application. The same applies to any future release. If the application is recompiled and linked against a later release, then the application must load an IBM MQ library that matches the later release, even if it continues to connect to a Version 7.1 queue manager.

Your application might not be linked to an IBM MQ library, but instead calls the operating system directly to load an IBM MQ library. If the library that is loaded is from Version 7.1 or later, IBM MQ checks the library is from the installation that is associated with the queue manager. If it is not, IBM MQ loads the correct library.

## Special migration considerations involving loading IBM MQ libraries

You might have been asked to modify the installation of an earlier IBM MQ release to satisfy the requirements of a build environment, or the IT standards in your organization. If you copied IBM MQ libraries to other directories, or created symbolic links, you ended up with an unsupported configuration. The requirement to move IBM MQ libraries to other directories was one of the reasons for changing the installation of IBM MQ on UNIX and Linux. You can now install IBM MQ into a directory of your choosing. You can also load IBM MQ libraries from the `/usr/lib` directory, which is normally on the default load path on UNIX and Linux systems.

A common IT standard or build environment requirement is to include IBM MQ libraries in the default load path on UNIX and Linux systems. IBM WebSphere MQ Version 7.1 has a solution. In Version 8.0 you can install IBM MQ into a directory of your own choosing, and IBM MQ can create symbolic links in `/usr` and its subdirectories. If you make a Version 7.1, or later, installation primary by using the **setmqinst** command, IBM MQ inserts symbolic links to the IBM MQ libraries into `/usr/lib`. As a result, the operating system finds the IBM MQ libraries in the default load path, if that includes `/usr/lib`.

Because IBM WebSphere MQ Version 7.1, or later, libraries transfer calls to the correct installation, defining that version installation as primary, also results in the correct libraries being loaded for any application that is built with a link to `/usr/lib`, regardless of which queue manager it connects to. Unfortunately, this solution does not work if you have a Version 7.0.1 installation on the server, because then you cannot define a Version 7.1, or later, installation as primary, and the Version 7.0.1 libraries do not load libraries from other installations. As an alternative to setting the later version installation primary, use **setmqenv** with the `-k` or `-l` options to achieve a similar result.

You can find more information in *Connecting applications in a multiple installation environment*.

## Command association

Examples of commands are **dspmqver**, **setmqinst**, **runmqsc**, and **strmqm**. The operating system must find a command in an IBM MQ installation. Many commands also require a queue manager as an argument and assume the default queue manager, if a queue manager name is not provided as a parameter.

Unlike loading libraries, if a command includes a queue manager as a parameter, the command is not switched to the installation that is associated with the queue manager. You must use the **setmqenv** command to set up your environment correctly, so that any commands that you issue are run from the correct installation. You can provide a queue manager as a parameter to **setmqenv**, to set up the command environment for that queue manager; see Figure 64 on page 464.

On Windows, the **setmqinst** command sets global environment variables, and **setmqenv** local environment variables, including the `PATH` variable to find commands.

On UNIX and Linux, the **setmqinst** command copies symbolic links for a subset of the commands into `/usr/bin`; see *External library and control command links to primary installation on UNIX and Linux*. The **setmqenv** command sets up local environment variables, including the search path to the binary folder in the installation directory.

**setmqenv** must be on the search path in order to run it. One reason for having a later version installation as primary is to be able to run **setmqenv** without having to configure the search path. If IBM WebSphere MQ Version 7.0.1 is installed on the server, no Version 7.1, or later, installation can be primary and IBM WebSphere MQ Version 7.0.1 does not have a **setmqenv** command. The consequence is, you must provide a path to run the **setmqenv** command to set up the command environment for any of the later version installations on the server.

Figure 64 on page 464 shows two examples of running **setmqenv** to set up the command environment for the copy of IBM MQ that is associated with the queue manager, QM1.

---

IBM WebSphere MQ for Windows Version 7.1  
" MQ\_INSTALLATION\_PATH\bin\setmqenv" -m QM1

IBM WebSphere MQ Version 7.1 for UNIX and Linux  
. MQ\_INSTALLATION\_PATH/bin/setmqenv -m QM1

---

Figure 64. Running setmqenv

**Related tasks:**

"Migrating IBM MQ library loading from an earlier version of the product to the latest version" on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

"UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version" on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** .

Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

"UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version" on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

"UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version" on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

"UNIX, Linux, and Windows: Staging maintenance fixes" on page 669

Use multiple installations of IBM MQ on the same server to control the release of maintenance fixes.

"UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version" on page 542  
Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.


"Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version" on page 540

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

**Related reference:**



“Coexistence” on page 581

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

#### **Related information:**

Changing the primary installation

Connecting applications in a multiple installation environment

 [Dynamic-Link Library Search Order](#)

External library and control command links to primary installation on UNIX and Linux

Features that can be used only with the primary installation on Windows

Installation configuration file, mqinst.ini

setmqenv

setmqinst

setmqm

**strmqm** -e CMDLEVEL

### **Application compatibility and interoperability with earlier versions of IBM MQ**

Connecting an application that is built against libraries shipped with a later version of IBM MQ to an earlier version IBM MQ is not supported. Avoid building applications against a later version, and redeploying them to a queue manager running at an earlier version, although some applications do work in practice.

IBM MQ applications do interoperate with applications running on earlier versions of IBM MQ, as long as they use no new function. IBM MQ clients can connect to queue managers running at an earlier version than the client, as long as the client uses no new functions.

An IBM MQ application that uses only functions provided by an earlier version of a queue manager can continue to send messages to the earlier version. It does not matter what version of IBM MQ an application is built on and connected to. It can exchange messages with an application connected to an earlier version of IBM MQ, as long as it does not use new function.

Consider these four cases; the first two cases are not supported though they might work in practice, the last two cases are supported. The first two cases require compatibility with an earlier version of IBM MQ. The last two cases rely on the interoperability between all versions of IBM MQ

1. Running an IBM MQ server application, built with a later version of IBM MQ, connecting to a queue manager running on a server with an earlier version of IBM MQ installed.
2. Running an IBM MQ client application, built with a later version of IBM MQ, on a client platform with an earlier client installation, connecting to a queue manager running on a server with a later version of IBM MQ installed.
3. Running an IBM MQ client application, built with a later version of IBM MQ, on a client platform with the later client installation, connecting to a queue manager running on a server with an earlier version of IBM MQ installed.
4. Exchanging messages between an IBM MQ client or server application, connected to a queue manager running on a server with a later version of IBM MQ installed, with applications connected to a queue manager running on a server with an earlier version of IBM MQ installed.

Plan to avoid the first two cases, as they are not guaranteed to work all the time. If you are running an incompatible configuration and you encounter a problem, you must rebuild your applications with the correct level of IBM MQ. You can then continue with problem diagnosis.

## Multi-installation and application loading

The new loading capability of IBM MQ libraries in Version 7.1, or later, does not relax the restriction that an application compiled and linked at a later release level must not directly load an IBM MQ library at an earlier release level. In practice the restriction is of less significance than in earlier releases, because as long as the operating system loads a library at the same or later level than the library the application was compiled and linked with, IBM MQ can call any other level of IBM MQ on the same server from Version 7.0.1 upwards.

For example, suppose you recompile and link an application that is to connect to a Version 7.0.1 queue manager using the libraries shipped with Version 7.1. At run time the operating system must load the Version 7.1 libraries for the application, even though the application connects to a Version 7.0.1 queue manager. IBM WebSphere MQ Version 7.1 detects the inconsistency and loads the Version 7.0.1 library for the application. The same applies to any future release. If the application is recompiled and linked against a later release, then the application must load an IBM MQ library that matches the later release, even if it continues to connect to a Version 7.1 queue manager.

### Examples

1. You decide to rebuild a client application. Can you deploy it to your production environment that contains some earlier versions of client and server platforms?

The answer is no, you must upgrade all the client workstations you deploy to, at least to the version of the client you have built. The queue managers running on earlier versions of IBM MQ do not have to be upgraded. In practice all the clients are likely to work, but for maintainability you must avoid running incompatible levels of an application and the IBM MQ client.

2. You deploy some IBM MQ queue managers at a new version level. You have an existing IBM MQ application that you use to send messages between the servers. Do you rebuild the application to deploy it onto the new servers? Can you deploy the old version onto the new servers?

The answer is, either. You can continue to deploy the existing version of the application onto all your servers, or you can deploy the rebuilt application onto the new servers. Either configuration works. IBM MQ supports running the existing application on later servers and sending messages from later application versions to earlier ones. What you must not do is to rebuild the application on the later version and redeploy it onto both the earlier and newer servers. IBM MQ does not support compatibility with earlier versions.

z/OS

### z/OS application stubs

The stub modules that are listed are link-edited with applications and exits. The version 7 stub modules might not work with Version 6.

- CSQASTUB
- CSQBRSSI
- CSQBRSTB
- CSQBSTUB
- CSQCSTUB
- CSQQSTUB
- CSQXSTUB

To take advantage of the new APIs introduced in Version 7.0, for example MQSUB and MQCB, you must link-edit your application with stubs shipped with Version 7.0 or later, for an LE program, the sidedeck, see Building z/OS batch applications using Language Environment . Such an application will not run on an earlier version of the queue manager.

## Application compatibility and interoperability with later versions of IBM MQ

IBM MQ applications run against later versions of a queue manager without recoding, recompiling, or relinking. You can connect an application that is built against libraries shipped with an earlier version of IBM MQ to a queue manager running at a later version of IBM MQ.

If you upgrade a queue manager to a later version, existing applications built against its earlier version work without change. Exceptions are noted in “Changes that affect migration” on page 580. Likewise applications connected to the IBM MQ Client, run against later versions of the client without recoding, recompiling, or relinking. You can deploy client applications built against earlier versions of the IBM MQ Client libraries to connect using later versions of the libraries.

All the following four cases are supported. The first two cases rely on the compatibility of later version of IBM MQ with applications built against earlier versions. The last two cases rely on the interoperability between all versions of IBM MQ.

You might change the operating environment as a prerequisite of migrating to a new level of queue manager. The operating environment changes, rather than changes in IBM MQ itself, might require application change, recompilation, or relinking. Sometime the operating environment change affects only the development environment, and the operating environment supports applications built at an earlier level. In which case, you might be able to run existing applications built at the older level of the operating environment. You might not be able to build any new applications until the operating environment is upgraded.

In the future, after you have migrated queue managers and clients to the latest release level, consider changing your applications to take advantage of new capabilities.

### z/OS

#### z/OS application stubs

The stub modules that are listed are link-edited with applications and exits. The Version 6.0 stub modules continue to work with Version 8.0.

- CSQASTUB
- CSQBRSSI
- CSQBRSTB
- CSQBSTUB
- CSQCSTUB
- CSQQSTUB
- CSQXSTUB

To take advantage of the new APIs introduced in Version 7.0, for example MQSUB and MQCB, you must link-edit your application with stubs shipped with Version 7.0 or later, for an LE program, the sidedeck, see Building z/OS batch applications using Language Environment . Such an application will not run on an earlier version of the queue manager.

## Compatibility between different versions of an IBM MQ MQI client and a queue manager

Any version and release of an IBM MQ MQI client can connect to any version and release of an IBM MQ queue manager. The MQI channel is automatically configured to the latest version that both the client and server support. If the client and server are different versions, the client application must use only the functions in the earlier version.

The compatibility between clients and queue managers applies only to the version and release ( V.R) of the product. The statement of compatibility does not necessarily apply to the modification and fix pack level ( M.F) of the product.

If there are known problems at a specific V.R.M.F of the product, an upgrade to a more recent fix pack for the same Version.Release is necessary.

When you upgrade a queue manager to a different version, you automatically upgrade IBM MQ libraries. The libraries are used by IBM MQ MQI client and server applications running on the same server as the queue manager. To access new functions from remote clients, you must also upgrade the IBM MQ MQI client installation on remote workstations. The IBM MQ MQI client includes the IBM MQ MQI client libraries.

Remote clients that have not been upgraded continue to work with an upgraded queue manager. The behavior of the client application might, in rare cases change. You must consult “Changes that affect migration” on page 580, to find out whether changes in the current version affect your client applications.

Remote clients that are connected to upgraded queue managers can use the new functions in the release. If an upgraded remote client is connected to a queue manager that has not been upgraded, it must not use new functions. In rare cases, the behavior of the client might change; see “Changes that affect migration” on page 580.

You can generally assume that upgrading the IBM MQ MQI client does not require you to recompile or relink the client application. You can also continue to use the same connection to the queue manager. If changes are required, they are identified in “Migrating a queue manager to the latest release” on page 512, for the particular migration path and platform you are concerned with.

The Client Channel Definition Table (CCDT) is an interface to customize the connection between an IBM MQ Client and a queue manager. Entries in the tables are client connections, which are defined using a queue manager. The version of a CCDT is the version of the queue manager used to define the client connections. If an IBM MQ MQI client uses CCDT to connect to a queue manager, the CCDT can be at a version greater than, less than, or equal to that of the client.

You can connect to a queue manager with an earlier IBM MQ Client or an earlier CCDT. If you are using a CCDT, and you plan to use new client channel configuration options, such as shared conversations, you must upgrade the CCDT, and therefore the IBM MQ MQI client installation to the new version.

## **MQI client: Client Channel Definition Table (CCDT):**

The client channel definition table has changed from Version 6.0 to Version 7.5. Existing clients must continue to use existing CCDTs. To use a Version 7.5 CCDT, you must update the client.

You can connect an IBM MQ MQI client application to any level of queue manager. If the client uses a CCDT, it must use a CCDT built by the same or earlier version of queue manager. If the client connects without using a CCDT, this restriction does not apply.

In a common migration scenario, if you upgrade a Version 6.0 queue manager to Version 7.5, and you do not create new CCDTs for its clients, the clients connect to the Version 7.5 queue manager without any changes being required. Client behavior might change as a result of changes to the queue manager.

Another common migration scenario is to update some queue managers and some clients to Version 7.5, leaving some queue managers and clients at Version 6.0. In this scenario, you want to update the CCDT for Version 7.5 IBM MQ MQI client connected to Version 7.5 queue managers to Version 7.5, so that the clients can exploit Version 7.5 function fully. The new clients also connect to Version 6.0 queue managers. Existing clients connect to both Version 6.0 and Version 7.5 queue managers. In order that Version 7.5 clients can exploit new Version 7.5 function, you must deploy a Version 7.5 CCDT to Version 7.5 clients. Version 6.0 clients must continue to use the Version 6.0 CCDT. Both sets of clients can connect to both sets of queue managers, regardless of the CCDT they are using.

If the client is an IBM MQ MQI client, the version of the IBM MQ MQI client libraries linked to by the client must be the same or greater than the version of the queue manager that was used to build the CCDT. If the client is a Java or JMS client, then the client must be built with versions of the IBM MQ JAR files that are the same or greater than the queue manager that was used to build the CCDT.

To upgrade a Version 6.0 IBM MQ MQI client to use a Version 7.5 CCDT, you must upgrade the IBM MQ MQI client installation to Version 7.5. Unless you decide to do so for other reasons, do not rebuild the client application.

To upgrade a Version 6.0 Java or JMS client to use a Version 7.5 CCDT, redeploy the IBM MQ JAR files to the client workstation. You do not need to rebuild the Java or JMS client with the new JAR files.

## **MQI client: Client configuration stanzas moved into a new configuration file:**

Client configuration information is moved from existing configuration stanzas into a new configuration file, `mqclient.ini`.

Moving client configuration information affects existing settings; for example:

- Set the TCP KeepAlive attribute for client connections in `mqclient.ini` ; for example:

```
TCP:
KeepAlive = Yes
```

An existing setting in `qm.ini` is ignored.

- Set ClientExitPath in `mqclient.ini` ; for example:

```
ClientExitPath:
ExitsDefaultPath=/var/mqm/exits
ExitsDefaultPath64=/var/mqm/exits64
```

An existing setting in `mqs.ini` is moved to the client configuration file when you upgrade the client. If you add values to `mqs.ini`, they are ignored.

- Set JavaExitsClasspath in `mqclient.ini`.

Do not continue to use the Java system property `com.ibm.mq.exitClasspath`. Existing settings continue to work, but they are deprecated. The setting in `mqclient.ini` has precedence over the Java system property.

**Related information:**

The IBM MQ classes for JMS configuration file  
Assigning channel exits for IBM MQ classes for JMS  
IBM MQ client configuration file

**MQI client: Default behavior of client-connection and server-connection channels:**

In Version 7.0 the default settings for client and server connection channels changed to use shared conversations. This change affects the behavior of heartbeats and channels exits, and can have an impact on performance.

Before Version 7.0, each conversation was allocated to a different channel instance. From Version 7.0, the default for client and server connections is to share an MQI channel. You use the **SHARECNV** (sharing conversations) parameter to specify the maximum number of conversations that can be shared over a particular TCP/IP client channel instance. The possible values are as follows:

**SHARECNV(0)**

This value specifies no sharing of conversations over a TCP/IP socket. The channel instance behaves exactly as if it was a Version 6.0 server or client connection channel, and you do not get the extra features such as bi-directional heartbeats that are available when you set **SHARECNV** to 1 or greater. Only use a value of 0 if you have existing client applications that do not run correctly when you set **SHARECNV** to 1 or greater.

**SHARECNV(1)**

This value specifies no sharing of conversations over a TCP/IP socket. Performance on distributed servers is similar to that for a value of 0. Client heartbeating (whether in an MQGET call or not) and read ahead are available, and channel quiescing is more controllable. You can usually use this setting with existing Version 6.0 client applications.

**SHARECNV(2) to SHARECNV(999999999)**

Each of these values specifies the number of shared conversations. If the client-connection **SHARECNV** value does not match the server-connection **SHARECNV** value, then the lowest value is used. The default value is **SHARECNV(10)**, which specifies 10 threads to run up to 10 client conversations per channel instance. However, on distributed servers there are performance issues with **SHARECNV** channels that can be eased by using **SHARECNV(1)** wherever possible.

For all **SHARECNV** values of 1 or greater, the channel supports the following features:

- Bi-directional heartbeats
- Administrator stop-quiesce
- Read-ahead
- Asynchronous-consume by client applications

You can also set the MQCONNX option, `MQCNO_NO_CONV_SHARING` and connect the application to a channel with **SHARECNV** set to a value greater than 1. The result is the same as connecting the application to a channel with **SHARECNV** set to 1.

**Performance**

The change in Version 7.0 to use shared conversations, and further enhancements introduced in Version 8.0, can impact performance on distributed servers. See Tuning client and server connection channels.

## Heartbeats

From Version 7.0, heartbeats can flow across the channel at any time in either direction. The SHARECNV(0) and Version 6.0 behavior is for heartbeats to flow only when an MQGET call is waiting.

## Channel exits

The behavior of a client or server connection channel exit changes when the channel is sharing conversations (that is, when you set **SHARECNV** to a value greater than 1). It is unlikely, but possible, that the change affects the behavior of existing exits. The change is as follows:

- Send or receive exits can alter the MQCD structure on an MQXR\_INIT call. The effect of these exits differs, depending on whether the conversation is shared with other conversations on the same channel:
  - If the MQCXP SharingConversations field passed to the exit instance is set to FALSE, this exit instance is the first, or only, conversation on the channel instance. No other exit can be changing the MQCD at the same time, and changes that are made to the MQCD can affect the way that the channel runs.
  - If the MQCXP SharingConversations field passed to the exit instance is set to TRUE, this exit instance is a subsequent conversation. It is sharing the channel instance with other conversations. Changes made to the MQCD in the exit instance are retained in the MQCD but do not affect the way the channel runs.
- Send, receive, and security exit instances can alter the MQCD, when the MQCXP SharingConversations field is set to TRUE. Exit instances on other conversations might be changing the MQCD at the same time. Updates written by one exit instance can be overwritten by another instance. It might be necessary to serialize access to the MQCD across these different exit instances to maintain the consistency of fields in MQCD.

Updating MQCD when the SharingConversations field is set to TRUE does not affect the way the channel runs. Only alterations made when the MQCXP SharingConversations field is set to FALSE, on an MQXR\_INIT call, change channel behavior.

### Related information:

Using sharing conversations

Channel-exit programs for MQI channels

Using read ahead

Stopping MQI channels

Tuning client and server connection channels

HeartbeatInterval (MQLONG)

SharingConversations (MQLONG)

ALTER CHANNEL

The Asynchronous consume sample program

### MQI client: MQPUT1 sync point behavior change:

An MQPUT1 call by an IBM MQ MQI client application that failed in IBM WebSphere MQ Version 6.0 can now sometimes succeed. The failure is returned to the application later, if it calls MQCMIT. For the change in behavior to occur, the MQPUT1 must be in sync point.

In the scenario, “Example call sequence that demonstrates the change in behavior” on page 472, an MQPUT1 call can succeed where it failed in Version 6.0. The result occurs when all the following conditions are met:

- Both client and queue manager are later than Version 6.0.
- The application program is connected to the queue manager as a client application
- MQPMO\_SYNCPOINT is set in the Put Message Options structure, MQPMO

You can make the IBM MQ MQI client behave like Version 6.0. Set **Put1DefaultAlwaysSync** to YES in the **CHANNELS** stanza of the client configuration file; see Figure 65.

---

```
Channels:
Put1DefaultAlwaysSync=YES
```

---

Figure 65. Add **Put1DefaultAlwaysSync** to *mqclient.ini*

### Example call sequence that demonstrates the change in behavior

---

1. MQCONN to queue manager from an IBM MQ MQI client application.
  2. MQPUT1 to a nonexistent queue with the MQPMO\_SYNCPOINT option
  3. MQDISC
- 

In IBM WebSphere MQ Version 6.0 the MQPUT1 call ends with MQCC\_FAILED and MQRC\_UNKNOWN\_OBJECT\_NAME(2085). Running with a client and server later than Version 6.0, the MQPUT1 call ends with MQRC\_NONE and MQCC\_OK.

#### Related information:

CHANNELS stanza of the client configuration file

## Queue manager migration

After upgrading an installation, queue manager migration might be required. Migration takes place when you start a queue manager.

**z/OS** On z/OS, queue manager migration is required after upgrading to a different version, release, or maintenance level. The upgrade changes the command level. The current command, or VRM, level is shown in the z/OS console log.


On other platforms, queue manager migration is always required for changes in the first two digits of the VRMF code. Changes in the maintenance and fix level, M and F in the VRMF code, never cause automatic queue manager migration. No migration was required for the upgrade from Version 7.0 to Version 7.0.1. The change from Version 7.0 to Version 7.0.1 did change the command level from 700 to 701. From Version 7.1 onwards, a change in the command level always requires queue manager migration, but if the change is shipped in a maintenance or fix pack, you have the choice of whether to increase the command level, and cause queue manager migration.

Command level always increases with a change in version or release. If you decide to use new function introduced in a maintenance level upgrade, you must change the command level. The converse is not the case. You do not have to change the command level when the fix level changes. You can decide to install the fix pack, but not use the new function. Whether or not you use the new function, the installation of the fix pack increases the maximum command level supported by the installation. Run the **dspmver** command to display the current maximum supported command level.

Queue manager migration is the process of converting persistent queue manager data from one version to another. Persistent queue manager data includes log files and data in the queue manager directory. The data records changes to objects such as messages, subscriptions, publications, queue managers, channels, queues, and topics.

Queue manager migration is required and largely automatic. **z/OS** On z/OS, you must migrate queue managers manually between compatibility mode and new function mode by setting the **OPMODE** parameter.



After migrating to a new release , or on z/OS, after setting the **OPMODE** to NEWFUNC, the queue manager cannot be restored to an earlier release level.

You can reduce the downtime and risk caused by queue manager migration, by verifying the new version first, using a different queue manager. Unless the platform supports queue manager coexistence, you need to perform the verification on a different server, or in a virtualized environment on the same server. If the platform you are upgrading supports queue manager coexistence, you can install the new version of IBM MQ on the same server, verify it, and minimize downtime to the time required to stop, backup, and restart the queue manager.

**Note:** If you are migrating a queue manager through multiple release levels, one level at a time, you must start the queue manager after each upgrade to migrate it. You must also start all the channels, to ensure they are migrated.

#### Related concepts:

“The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423  
IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

 “The version naming scheme for IBM MQ for z/OS” on page 422

On IBM MQ for z/OS, releases have a three-digit Version, Release, and Maintenance (VRM) level code. The code is significant; it identifies the service life of a release. To run a queue manager at a different VRM level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

“Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )” on page 453

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

 “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

#### Related tasks:

“Migrating a queue manager to the latest release” on page 512



The procedures for migrating a queue manager to the latest release are detailed in the following topics.

#### Related information:

dspmqver

## Reverting a queue manager to a previous version

You can remove an upgrade before you have started a queue manager. After a queue manager has been started, if you remove the upgrade, the queue manager will not work.

  You must back up your system before starting migration. You can either back up queue manager data, or use a backup queue manager; see Backing up and restoring IBM MQ. To back up, you must stop the queue manager.

You can reduce the downtime and risk caused by queue manager migration, by verifying the new version first, using a different queue manager. Unless the platform supports queue manager coexistence, you need to perform the verification on a different server, or in a virtualized environment on the same server. If the platform you are upgrading supports queue manager coexistence, you can install the new version of IBM MQ on the same server, verify it, and minimize downtime to the time required to stop, backup, and restart the queue manager.

▶ **z/OS** On z/OS, before Version 7.0.1, it is possible to revert to a previous level, as long as you have applied the correct PTFs.

▶ **z/OS** On z/OS, from Version 7.0.1 onwards, it is impossible to revert to an earlier release after running with **OPMODE NEWFUNC**. Otherwise, you can backward migrate, as described in Migration PTFs.

▶ **z/OS** On z/OS, you must migrate queue managers manually between compatibility mode and new function mode by setting the **OPMODE** parameter. If you have never switched a queue manager to new function mode, you can still run it against the earliest release it is compatible with. You must have applied compatibility PTFs to the earlier release before starting a queue manager at the new command level. The compatibility level is shown in the log.

#### Related concepts:

“Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS)” on page 453  
You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

▶ **z/OS** “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452  
You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

#### Related reference:

▶ **z/OS** “z/OS: OPMODE” on page 677  
The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. To access V8.0 capabilities, change the value of **OPMODE** to **OPMODE=(NEWFUNC,800)**. To restrict the use of new capabilities, and retain the ability to revert the queue manager to its earlier level, leave **OPMODE** at its default setting, **OPMODE=(COMPAT,800)**.

#### Related information:

Backing up and restoring IBM MQ

## IBM MQ MQI client migration

IBM MQ MQI client migration is the process of converting IBM MQ MQI client configurations, and client and server channels from one version to another. Client migration can take place after upgrading the IBM MQ MQI client, and are reversible.

Client migration on the client workstation is optional and manual. Client migration on the server is required and automatic. See “Changes that affect migration” on page 580 for a list of any client changes. You must upgrade an IBM MQ MQI client before migrating a client workstation to make use of new configuration options. You can make configuration changes to client and server connection channels on the server, but they have no effect on a client workstation, until the client is upgraded.

An example of client migration performed at the client workstation is to manually migrate configuration settings to the `mqclient.ini` configuration file.

An example of combined client and server migration is the deployment of a new client connection definition table (CCDT). To use a new version of the CCDT, generate the table on a queue manager that is at the new code level. Deploy the table to clients that are going to use it. To deploy the table to a client, you first must update the client to at least the same level as the queue manager that created the table.

IBM MQ MQI clients are interoperable with earlier and later versions of IBM MQ. Upgrading the IBM MQ MQI client makes new function available to client applications, and is important to maintain the service level. Migrating an IBM MQ MQI client gives it access to new configuration options.

The IBM MQ MQI client libraries, such as `mqic.dll`, are dynamic, and the application linkages to the libraries do not normally change. You do not relink a client application to pick up new IBM MQ client libraries. The client picks up the new library next time the library is loaded by the client application. Do not move libraries from their installed directory. Linking to libraries in anything other than their installed directory is an unsupported configuration.

**Related concepts:**

“Application compatibility and interoperability with earlier versions of IBM MQ” on page 465  
Connecting an application that is built against libraries shipped with a later version of IBM MQ to an earlier version IBM MQ is not supported. Avoid building applications against a later version, and redeploying them to a queue manager running at an earlier version, although some applications do work in practice.

“Application compatibility and interoperability with later versions of IBM MQ” on page 467  
IBM MQ applications run against later versions of a queue manager without recoding, recompiling, or relinking. You can connect an application that is built against libraries shipped with an earlier version of IBM MQ to a queue manager running at a later version of IBM MQ.

**Related reference:**

“MQI client: Client Channel Definition Table (CCDT)” on page 469  
The client channel definition table has changed from Version 6.0 to Version 7.5. Existing clients must continue to use existing CCDTs. To use a Version 7.5 CCDT, you must update the client.

“MQI client: Default behavior of client-connection and server-connection channels” on page 470  
In Version 7.0 the default settings for client and server connection channels changed to use shared conversations. This change affects the behavior of heartbeats and channels exits, and can have an impact on performance.

“MQI client: MQPUT1 sync point behavior change” on page 471  
An MQPUT1 call by an IBM MQ MQI client application that failed in IBM WebSphere MQ Version 6.0 can now sometimes succeed. The failure is returned to the application later, if it calls MQCMIT. For the change in behavior to occur, the MQPUT1 must be in sync point.

## **Migrating an IBM MQ MQI client on UNIX platforms, and Windows to the latest version**

Before migrating an IBM MQ MQI client, create a migration plan. Stop all IBM MQ activity on the client workstation. Upgrade the IBM MQ MQI client installation. Make any essential configuration and application changes.

### **Before you begin**

1. Create a migration plan. Use the planning task, “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485, as a guide.

### **Procedure**

1. Review the latest IBM MQ requirements; see Checking requirements.
2. Review all the changes in IBM MQ that affect you; see What's changed in IBM MQ Version 8.0 and “Changes from WebSphere MQ Version 7.0.1 to IBM MQ Version 8.0” on page 584
3. End all IBM MQ activity on the workstation.
4. Upgrade the client.
  - To upgrade a client installation on a workstation; see the appropriate topic for the platform, or platforms, that your enterprise uses:
    - Client installation procedure on an AIX workstation
    - Client installation procedure on an HP-UX workstation

- Client installation procedure on an Linux workstation
- Client installation procedure on an Solaris workstation
- Client installation procedure on an Windows workstation
- To upgrade a client installation on a server; see Installing an IBM MQ MQI client on the same computer as the server

## What to do next

After upgrading the IBM MQ MQI client, you must check the client channel configuration, and verify that your IBM MQ MQI client applications work correctly with the latest version of the product.

## IBM i: Migrating an IBM MQ classes for JMS and Java client

If you have IBM MQ Java SupportPac MA88 installed, you must uninstall it first.

### Before you begin

#### SupportPac MQ88 is installed.

If you try to install the latest version of IBM MQ classes for Java anyway, the installation fails with a warning requesting you to uninstall the old client. You must follow the steps in this task to uninstall IBM MQ classes for Java and IBM MQ classes for JMS.

#### A previous version of IBM MQ classes for Java is installed.

Installation of the latest version of IBM MQ classes for Java uninstalls the previous version automatically. Do not follow the steps in this task.

### About this task

The steps in this task uninstall the IBM MQ classes for JMS and Java.

### Procedure

To uninstall the previous IBM MQ Java client:

1. Delete the QMQMJAVA library and the /QIBM/ProdData/mqm/java directory, by issuing the command:  
DLTLICPGM LICPGM(5648C60) OPTION(\*ALL)
  2. If the previous step failed to delete the IFS directory /QIBM/ProdData/mqm/java and its subdirectories, use the **EDTF** command, for example:  
EDTF STMF('/QIBM/ProdData/mqm')
- and select option 9 against the java directory.

## Restoring an IBM MQ MQI client and client connection to a previous release

If you restore an IBM MQ MQI client to a previous code level, you must undo the configuration changes manually.

### About this task

It is unusual to restore earlier IBM MQ MQI client libraries to a workstation. The principal tasks are listed in the following steps.


### Procedure

1. End all IBM MQ activity on the workstation.
2. Uninstall the IBM MQ MQI client code for the current version.
3. Follow the client installation procedure for the appropriate platform to install the IBM MQ MQI client for the previous release.
4. If you configured a Client Connection Definition Table (CCDT) for a queue manager using the current version, revert to using a table created by a queue manager for the previous release.

The CCDT must always be created by a queue manager on the same, or earlier, release to the client.

## Application migration and interoperation

IBM MQ supports running applications compiled and linked against previous versions of IBM MQ, with later levels of IBM MQ.

To migrate an application to run with a new level of IBM MQ, disconnect an application from the queue manager. Reconnect it when the queue manager is running again. However, it takes only one small difference in the interface between IBM MQ and the application to break an application, or make it behave wrongly. Sometimes a problem does not show up for a long time. For this reason, you must always test your applications against a new version of IBM MQ. The suggested extent of testing varies depending on the extent of the changes in IBM MQ ; see  "Characteristics of different types of upgrade on z/OS" on page 453 or "Characteristics of different types of upgrade" on page 454.

Application migration refers to four kinds of changes.

1. Application changes that are consequent to upgrading the operating environment along with the queue manager. Rarely, linkage conventions change. The most likely reason for a linkage change is switching from 32 bit to a 64 bit environment. If you are using SSL or TLS you might have to relink with a new secure library.
2. Changes that you must make to the application in order to run an application against a new level of queue manager. Changes of this sort are uncommon. However, you must check "Changes that affect migration" on page 580 to see if any changes might affect your applications.
3. Changes that are not required, but that you might want to make in future, perhaps because you have a business reason to modify an application.
4. Changes to applications that are supplied by IBM, or other vendors, that require you to run migration utilities. The utilities convert the applications to running on the new version of IBM MQ.

Do not load IBM MQ libraries from an earlier level. IBM MQ does not support connecting server applications loading libraries from the earlier level to connect to a later level of queue manager. On UNIX, Linux, and Windows platforms, the application load path must be set up to the location of the IBM MQ server libraries. You do not have to recompile and relink an application. Applications compiled and linked against an earlier version of IBM MQ can load libraries from a later version.

On IBM i, UNIX, Linux, and Windows, from Version 7.1 onwards, IBM MQ loads the library from the installation the application is connecting to. An application must initially load a library of at least the same level as the application linked to. IBM MQ then loads the correct version of the library from the

installation that the queue manager is associated with. If you have two installations of the same version, but at different fix levels, IBM MQ chooses which library to load. The choice is based on the queue manager the application is connected to. If an application is connected to multiple queue managers, it is possible that multiple libraries are loaded.

To help you write applications that can exchange messages with earlier versions of the product, IBM MQ provides data type versioning. Data type versioning assists you in exchanging messages that are compatible with target queue managers. A good programming practice is to set the version number of a data structure explicitly. Do not assume that the default version is the one you require. By setting the version explicitly, you are forced to look up what version to use. The description of the data type version tells you what level of queue manager supports that version.

It is poor practice to set the data type version to the current version. If you recompile your program against a new version of IBM MQ, the data type version might change with unexpected consequences.

Client applications are more likely to connect to different queue managers than applications written for a specific server. Plan carefully when writing an application that is to connect to different versions of a queue manager, and to queue managers on different platforms. The default values of some IBM MQ constants, such as MQPMO\_SYNCPOINT, MQPMO\_NO\_SYNCPOINT differ between platforms. Some functions are not available on all platforms.

You must be aware of, and code to, the capabilities of all the queue managers the application interacts with. It requires planning and design to write an application that works with different versions of a queue manager. There is no API provided with IBM MQ to restrict an application to a function subset common to the set of queue managers it interacts with. To improve interoperability, some developers choose to provide an MQI wrapper layer, or use MQI API exits, to control the functions programs use.

#### **Related concepts:**

“Application compatibility and interoperability with earlier versions of IBM MQ” on page 465  
Connecting an application that is built against libraries shipped with a later version of IBM MQ to an earlier version IBM MQ is not supported. Avoid building applications against a later version, and redeploying them to a queue manager running at an earlier version, although some applications do work in practice.

“Application compatibility and interoperability with later versions of IBM MQ” on page 467  
IBM MQ applications run against later versions of a queue manager without recoding, recompiling, or relinking. You can connect an application that is built against libraries shipped with an earlier version of IBM MQ to a queue manager running at a later version of IBM MQ.

## **Queue manager cluster migration**

You can migrate queue managers in a cluster all at once, or one at a time, which is called a staged migration. Migrate full repository queue managers in a cluster before partial repository queue managers.

Cluster queue managers can participate in clusters with other queue managers running at different versions, which is why a staged migration is possible. Being able to stage a migration is important, as migrating each queue manager in a cluster takes time. By staging the migration, which leaves other queue managers that are in the cluster running, you reduce the effect of queue manager downtime on applications.

Migrate queue managers with full repositories first. Then migrate the other queue managers, which have partial repositories, one at a time. Complete migration of the entire cluster before starting to use new functions.

If you do have to start using new functions before completing migration of the entire cluster, you might need to refresh the partial repositories. After each migration of a queue manager with a partial repository, issue the **REFRESH CLUSTER** command on the newly migrated queue manager. The command updates the cluster records in the newly migrated queue manager, potentially receiving updates for any new

attributes. Do not do this step if you migrated the entire cluster before using new function. The **REFRESH CLUSTER** command takes a long time for all the changes to work through the cluster.

**Note:** For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.

If full repositories are not migrated before partial repositories, the cluster continues to work, but without all the new features in a version working as expected. To work predictably, the full repository queue managers must be at the new command level to be able to store information from the rest of the cluster that arises from using new features.

For example, the information might be a new channel attribute, such as shared conversations, which were introduced in Version 7.0. Information about the shared conversation attribute of a channel between two other Version 7.0.1 queue managers can be stored in a Version 7.0 full repository, but not in a Version 6.0 repository. If information about a channel with the shared conversation attribute is updated from the Version 6.0 full repository, the definition loses its shared conversation attribute. “How mixed version cluster repositories are updated” explains how information is updated in a mixed-version cluster.

**Note:** If a queue manager is a member of a cluster, and is running at a release earlier than Version 6.0, you must migrate the queue manager to Version 7.0.1, before migrating it to the latest release. You must start the queue manager after the first migration step, before proceeding to Version 8.0.

### **How mixed version cluster repositories are updated**

Repositories store records for an object in a cluster in the version of the record format that matches the version of the queue manager hosting the repository. Repository queue managers forward object records, before they are stored, in the format that they are received in. The recipient ignores fields from a newer version, and uses default values for fields that are not present in the record.

Cluster repositories hold records that represent objects, for example, a queue record represents a cluster queue. A full repository holds records for all objects in the cluster. Partial repositories hold records for local objects and remote objects that are used locally. A repository record can hold information only about attributes at the same command level as the queue manager holding that repository. So for example, a Version 6.0 repository contains only Version 6.0 level attribute information. A Version 8.0 repository contains all Version 6.0 records, plus Version 8.0 records containing additional Version 8.0 attributes.

A repository stores a record it receives in its own version. If the record it receives is at a later version, the later version attributes are discarded when the record is stored. A Version 6.0 queue manager receiving information about a Version 8.0 queue manager stores only Version 6.0 information. A Version 8.0 repository receiving a Version 6 record stores default values for attributes introduced in the version 7. The defaults define the values for the attributes that are not included in the record it receives.

A repository normally sends records in its own version format, which is the same as the format it has stored them in. There is one exception to this rule. When a full repository receives a record from a partial repository, it is immediately forwarded in the same format. So if a Version 6.0 full repository were to receive a record from a Version 8.0 partial repository, it would forward the Version 8.0 record. It sends the record to any other full repositories, and any other partial repositories that have subscriptions that match the record.

A partial repository reflects whichever full repository sent it the latest update to a record. As a consequence, you might see the information held by a Version 8.0 partial repository for new Version 8.0 attributes changing unexpectedly. The values might change from actual Version 8.0 information to default values. The changes occur if the full repositories in the cluster are at different levels. Migrate full repositories first to avoid instability.

A partial repository sends information about its objects to a full repository periodically at least once every 27 days. Information is sent about any object when it is altered or defined. See [How long do the queue manager repositories retain information?](#)

After migrating all full repositories to Version 8.0, some attributes might hold default values. The attributes might hold default values in place of actual values, if a repository has not received an update. You can refresh the repository in either of two ways:

- Alter the object which the record containing default values represents, for example, using ALTER QL for a local queue. The alteration forces the local repository to send the record again.
- Issue the **REFRESH CLUSTER** command on the partial repository which holds the record containing default values. **REFRESH CLUSTER** forces the partial repository to discard the record containing default values and get a new record as required.

**Note:** For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See [Refreshing in a large cluster can affect performance and availability of the cluster.](#)

In summary, for the most predictable, and fastest migration, when you stage cluster migration do these steps in the following order:

1. Migrate the queue managers with full repositories.
2. Migrate the queue managers with partial repositories.
3. Start using new function in the cluster.

**Related information:**

[How long do the queue manager repositories retain information?](#)

## Queue-sharing group migration

▶ z/OS

You can combine queue managers from different releases in a queue-sharing group. Limit the time you manage a mixed group to only as long as it takes to migrate all the queue managers to the same command level. You cannot combine a queue manager at Version 8.0 or later in the same queue-sharing group as queue managers earlier than Version 7.0.1. You must update all queue managers in a queue-sharing group with a coexistence PTF, before migrating any of them. All queue managers in the queue-sharing group must be at the same version before any queue manager can be migrated to Version 8.0.

When you migrate queue managers in a queue-sharing group, aim to migrate all the queue managers to the new version as soon as you can. Queue-sharing groups can contain queue managers with a restricted mixture of versions. A mixture of queue managers in a queue-sharing group is supported so that you can migrate and test the upgrade of each queue manager. Migrate each queue manager, one at a time, leaving the queue-sharing group running. Mixed groups are harder to administer, than if all the queue managers are at the same version.



### Related reference:

“z/OS: MQSC commands in a mixed queue-sharing group” on page 584

Existing **MQSC** commands using new keywords and attribute values can be entered for routing to a migrated queue manager. You can enter the commands on any queue manager. Route the commands using **CMDSCOPE**. Commands with new keywords and attribute values, or new commands, routed to a previous version of queue manager, fail.

“z/OS: Properties of objects in a mixed queue-sharing group” on page 584

Attributes that did not exist in earlier versions can be created and altered on queue managers of a later version in a mixed queue-sharing group. The attributes are not available to queue managers in the group that are at an earlier level.

“z/OS: Queue sharing group coexistence” on page 583

A queue-sharing group can contain queue managers running on IBM WebSphere MQ Version 7.0.1, and on later releases. The queue managers can access the same shared queues and other shared objects. Queue managers running earlier versions of the product must have the coexistence PTF applied for the latest release.

## Migrating a queue manager in a high-availability configuration

Follow standard procedures to migrate a queue manager that is part of a high-availability configuration. (On platforms other than z/OS.)

High-availability configurations of queue managers can increase the availability of IBM MQ applications. If a queue manager, or server fails, it is restarted automatically on another server. You can arrange for IBM MQ MQI client applications to automatically reconnect to the queue manager. Server applications can be configured to start when the queue manager starts.

**distributed** **IBM i** High-availability configurations are implemented either by using a high-availability cluster solution or by using multi-instance queue managers. Red Hat Cluster Suite or Microsoft Cluster Service (MSCS) are examples of high-availability cluster solutions.

**z/OS** On z/OS there are several alternative techniques to increase queue manager availability; see Availability. Migration considerations on z/OS depend on the availability techniques that are employed, and are not described in this topic. The term high-availability configuration refers only to queue managers in configurations on platforms other than z/OS.

The overall principles involved in queue manager migration in a high availability configuration are the same, whether you are migrating a multi-instance queue manager or a high-availability cluster. In either case, the principles are as follows:

1. You must not restart a queue manager at a lower command level than the one it was previously running.
2. You cannot upgrade the code an active queue manager is running.
3. You cannot back up an active queue manager.

## Overall steps to migrate a queue manager in a multi-instance queue manager configuration

The following terms are relevant:

### active queue manager instance

A queue manager instance that has been started permitting standby instances, and is running.

### standby queue manager instance

A queue manager instance that has been started permitting standby instances, and is in standby. It is ready to take over from the active instance automatically.

Base your migration procedure on the following steps.

1. Before you start the migration process, create a different queue manager on a server on which you have installed the upgrade. Test the upgrade by performing whatever verification checks that your organization requires.
2. If you have a pool of servers that you pick from when starting a queue manager instance, upgrade IBM MQ on the servers that are in the pool and are neither active or acting as a standby.
3. Stop the standby queue manager instance. Make sure that you have no system management procedure running that restarts the instance automatically.
4. If you do not have a pool of servers, upgrade IBM MQ on the server that was running the standby instance.
5. Decide whether downtime or recoverability is more important in the migration:  
Follow these steps if recoverability is more important, and you must take a backup:
  - a. Stop the active queue manager instance, without switching to any standby.
  - b. Back up the queue manager.
  - c. Start a queue manager instance, permitting standbys, on one of the upgraded servers.
  - d. If you have a pool of upgraded servers, start another one, permitting standbys.  
If availability is more important, follow this procedure; you do not take a backup.
  - a. Start a queue manager instance as a standby on one of the upgraded servers.
  - b. Stop the active queue manager instance, switching to the standby.
  - c. If you have a pool of upgraded servers, start another one, permitting standbys.
6. Upgrade the IBM MQ code on the server that was the active queue manager instance, and start it as the standby instance if you have not already started a standby.

## Overall steps to migrate a queue manager in a high-availability cluster

The following terms are relevant:

### **active server**

The running server or active queue manager instance

### **passive server**

A server that is ready to take over from the active server automatically.

### **inactive server**

A server that is not prepared to take over automatically. The server might have been removed from the cluster, or taken offline in some way.

Base your migration procedure on the following steps. The details depend on the specific commands in the cluster concerned.

1. Before you start the migration process, create a different queue manager on a server on which you have installed the upgrade. Test the upgrade by performing whatever verification checks that your organization requires.
2. If you have four servers available, you can form two cluster pairs.  
With two pairs, the queue manager can continue to run in a cluster-pair at the old command level. When you are ready, you can transfer the queue manager to the pair of servers at the new command level.
3. Remove a passive server from the cluster. Make sure that the cluster cannot automatically restart the server. The server is made inactive.
4. If a high-availability cluster is using a common location for IBM MQ code, you must create a second location for the upgraded code.
5. Install, or upgrade, IBM MQ code using the server that is not now running the queue manager.

6. Verify the upgrade by creating a different queue manager on the server, and performing whatever verification checks that your organization requires.
7. If more than half the servers remain in the cluster, remove a server, upgrade IBM MQ, and verify the upgrade. Each server is made inactive as part of the process. Continue until half the servers are upgraded.
8. If your active server is part of a remaining cluster, deactivate the passive servers so that the cluster cannot reactivate them automatically.
9. Decide whether downtime or recoverability is more important in the migration:  
Follow these steps if recoverability is more important:
  - a. Stop the queue manager and remove the server from the cluster.
  - b. Back up the queue manager.Or this step, if downtime is more important:
  - a. Add the migrated servers back into the cluster, as passive servers. Then switch the remaining server in the high-availability server cluster over to one of the passive servers. The switch causes the running queue manager to stop, restarts it on one of the passive servers.
10. Upgrade any remaining high-availability servers, and add them back into the cluster.

**Related tasks:**

“Windows: Migrating an MSCS configuration” on page 573

Migrate queue managers in MSCS configuration one node at a time, following these instructions.

---

## Introduction to changes for Windows on IBM MQ Version 8.0

A number of changes have been made for IBM MQ Version 8.0 for Windows. You must understand these changes before planning any migration tasks using IBM MQ Version 8.0 for Windows.

### Installing a single copy of the product

If you have an existing previous version of the product on your system, and want to upgrade to the latest version, you have various options. You can either:

- Uninstall the previous version and then install the latest version,
- Install the new copy alongside the currently installed one and uninstall the original at a later time. See “Installing the product alongside an existing version,” or
- Perform a migration install, electing to replace the currently installed version when prompted.

After you have installed the product, start each queue manager and its data migration takes place. This includes the migration of queue managers from 32 bit to 64 bit.

### Installing the product alongside an existing version

If you want to install another version of the product alongside your existing product you can do so. See “Multiple IBM MQ installations” on page 436 and “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500 for further information.

When you install the new version of the product, run the setmqm command to associate the queue managers with the new installation.

Start each queue manager in turn and its data migration takes place.

## Upgrading one of a pair (or more) of installations

If you already have, for example, a IBM WebSphere MQ Version 7.5 installation and an IBM MQ Version 8.0 installation on a machine, upgrading the Version 7.5 installation to Version 8.0 requires the following additional step.

When you start the Version 8.0 installer, you are asked whether you want to **Install a new instance** or **Maintain or upgrade an existing instance**.

However, only the other Version 8.0 installation, or installations, are displayed; not the Version 7.5 installation in the selection box. At this point, select **Install a new instance**.

After the splash screen has been displayed, a second panel appears, which lists any older installations that you can upgrade to Version 8.0 using the Version 8.0 installer.

On this panel, select **Upgrade 7.5.0.n Installation 'Installation m'**, and then click **Next**.

## Change of digital signature algorithm

The IBM MQ programs and installation image are digitally signed on Windows to confirm that they are genuine and unmodified.

In older releases the product was signed using the SHA-1 with RSA algorithm.

In IBM MQ Version 8.0, the SHA-256 with RSA algorithm is now used. Some old versions of Windows do not support the new digital signature algorithm, but those versions are not supported by IBM MQ Version 8.0.

See Hardware and software requirements on Windows(tm) systems, and ensure that you install IBM MQ Version 8.0 on a supported version of Windows.

## Existing applications

All applications that were built with previous versions of the product continue to work with a 64 bit queue manager.

All applications using the C++ object interface need to be rebuilt; applications using the C interface are not affected.

## Exits

Queue manager exits on Windows 64-bit operating systems must be compiled as 64-bit exits. Any 32-bit queue manager exits must be recompiled before they can be used with a 64-bit queue manager. If you try to use a 32-bit exit with a 64-bit queue manager on IBM MQ Version 8.0, an AMQ9535 "invalid exit" error message is issued.

## Clients

32-bit client applications can connect transparently to queue managers from all supported versions of IBM MQ. This includes 64-bit IBM MQ Version 8.0.

## Samples

The samples for the C and C++ languages are compiled as 64-bit.

### Related information:

Windows: changes for IBM MQ Version 8.0

Directory structure on Windows systems

Hardware and software requirements on Windows systems

---

## IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i


Before migrating from one version to another, read this topic. Create your own migration plan based on the outline in the planning topic.

If there are concepts about migration you do not understand, read “Introduction to IBM MQ migration” on page 421 first.

### About this task

Use the following steps as a guide to creating a migration plan.

### Procedure

1. Review the IBM MQ system requirements for Version 8.0.  
See IBM MQ System Requirements.
2. Decide whether to run the current version of your product and the latest version on the same server.  
Note that if you require a multi version installation, and you are using a version of the product prior to IBM WebSphere MQ Version 7.0.1.6, you must install IBM WebSphere MQ Version 7.0.1.6 first.
3. Review all the changes in IBM MQ that affect you.  
See What's changed in IBM MQ Version 8.0.
4. Review performance changes.  
Performance reports are published as Supportpacs; see IBM MQ - SupportPacs by Product.
5. Review the latest README file for the product you are working with.  
See IBM MQ and MQSeries product READMEs.
6. Plan the sequence and timing of queue manager upgrades.  
If the queue manager is part of a queue manager cluster, you must migrate the queue managers that are full repositories first.  
If the queue manager is part of a high availability cluster, plan the migration to minimize downtime and maximize availability; see “Migrating a queue manager in a high-availability configuration” on page 481.
7. Plan to migrate your queue manager to the latest version.  
UNIX platforms - Migrating a queue manager from a previous release to the latest release.  
Windows - Migrating a queue manager from a previous release to the latest release  
 IBM i - Migrating a queue manager from a previous release to the latest release or Migrating a queue manager from a previous release to the latest release, alternative method  
Backing up queue manager data is part of the queue manager migration task.  
An alternative approach to backing up queue manager data, is to install and configure a new server. Test the latest version with a new queue manager on the new server. When you are ready to go into production on the latest version, copy the queue manager configuration and data, to the new server.
8. Plan to update any manual or automated procedures you have written with changes to messages and codes.
9. Decide on what regression tests to perform before putting the queue manager into production on the latest version.

- Include the procedures and applications you identified in steps 6 and 7 in your regression tests.
10. Plan to upgrade your IBM MQ MQI client installations to the latest version.
  11. Plan to upgrade your client and server applications to use new functions in the latest version.

## **z/OS: Migration planning to the latest release**

Create a migration plan for IBM MQ for z/OS to upgrade to the latest release.

### **About this task**

The goal of this task is for you to create your own plan to migrate your queue managers to the latest release. Incorporate the task to migrate a queue manager, “Migrating IBM MQ for z/OS - order of tasks” on page 547, into your plan.

Migration plan overview for your enterprise

| <b>Migration phase</b>                                                                                                                                                                      | <b>Required tasks</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Phase I, before migration. See “Preparing to migrate a single IBM MQ for z/OS queue manager” on page 554 for further information.                                                           | Preparing each queue manager in your enterprise for migration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Phase II, migrate each single queue manager in the order listed. See “Migrating a single IBM MQ z/OS queue manager to the next release of the product” on page 558 for further information. | Carry out this process for each queue manager in the order shown: <ol style="list-style-type: none"> <li>1. z/OS queue managers in a queue sharing group (QSG).</li> <li>2. z/OS queue managers not in a QSG (if any).</li> <li>3. Queue managers in a full repository on platforms other than z/OS. See “Migrating a queue manager to the latest release” on page 512 for further information.</li> <li>4. z/OS partial repository queue managers in a QSG.</li> <li>5. z/OS partial repository queue managers not in a QSG.</li> <li>6. z/OS queue managers not in a QSG or cluster.</li> <li>7. Partial repository queue managers on platforms other than z/OS.</li> <li>8. Queue managers not in a cluster on platforms other than z/OS.</li> </ol> |
| Phase III, post migration. See “Post migration tasks” on page 564 for further information.                                                                                                  | Carry out a full regression test, and then explore the new function available to you.<br><br>Optionally, at any time in the process, upgrade your client libraries if necessary, recompile your clients using the additional features provided by the new version, and deploy the clients.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

### **Procedure**

1. Review the IBM MQ system requirements for the latest version.  
See System requirements for IBM MQ.
2. Review all the changes in the product that affect you. See What's changed in IBM MQ Version 8.0 for further information.
3. Review performance changes.  
Performance reports are published as Supportpacs; see IBM MQ - SupportPacs by Product.

4. Review the backward and coexistence (or migration and toleration) PTFs for your current version of the product. See IBM MQ Support, Migration PTFs.

These PTFs must be applied to your current version of the product to enable you to revert your queue managers to the current version, after the queue managers have been started at the target version.

Note, that you can have different versions of queue managers coexisting in the same queue-sharing group.

If you are unsure which migration PTFs you require, run the following command SMP/E:

```
REPORT MISSINGFIX ZONES(mqgtzone) FIXCAT(IBM.Coexistence.MQ.V8R0M0)
```

See FIXCAT and IBM MQ Migration Installation for further information.

**Attention:** If a PTF requires a rebind of Db2 plans, the PTF is shipped with ++HOLD(ACTION), indicating the need for this process. In such a case, see Migrating Db2 tables to bind the plans before starting migration.

Other FIXCAT categories are listed in IBM Fix Category Values and Descriptions.

5. Plan to install the latest version early code, and activate for all queue managers on the LPAR. See Installing early code for more information.

Note that:

Before migration, all systems that are running queue managers that you plan to migrate to the latest version must have the early code for that version installed and running. Queue managers in queue-sharing groups that contain queue managers that are to be migrated, must also be running the early code.

A queue manager must use the early code from the same release level, or a later release level.

6. Consider using aliases for the IBM MQ libraries.

For example, use the IDCAMS utility with the DEFINE command:

```
DEFINE ALIAS(NAME(MQM.SCSQANLE)RELATE(MQM.V710.SCSQANLE))
```

You can use MQM.SCSQANLE, where applicable, in your STEPLIB, and it resolves to the actual data set. When you migrate to a new release, change the alias definition, rather than changing all the places in your JCL where the library is referenced.

This process has the most benefit for your server application programs, because you can get all of the programs to refer to the new libraries at the same time.

7. Plan the sequence and timing of queue manager upgrades.
  - You must install the backward and coexistence (or migration and toleration) PTF to bring the previous version queue managers up to the latest maintenance level for that version.
  - You must install the PTF on all members of a queue-sharing group, before you migrate any queue managers to the latest version. You can install the PTF one member at a time, leaving the other members running.
  - If the queue manager is a member of a queue manager cluster, you must consider the order of migration of queue managers in the cluster; see “Queue manager cluster migration” on page 478.
  - Check that any products that require the previous version of the product also support the new version.
8. Plan to update any manual or automated procedures you have written with changes to messages and codes.
9. Plan to update applications that might be affected by changes.

Update the IBM MQ library in the application STEPLIB concatenations to the latest version.

Consider whether the application must be able to run on both the previous version and the latest version. You might be able to change the application to be compatible with both code levels. If

you cannot, you can query the queue manager command level, and make the code conditional on the command level. Call MQINQ setting the MQIA\_COMMAND\_LEVEL selector.

10. Decide on what regression tests to perform before enabling the new function in the latest version.

The **OPMODE** parameter controls a staged migration from the previous version to the latest version.

Initially set **OPMODE** to (COMPAT,800) to make sure that you can go back to using an earlier version of the product. Once you are satisfied with the stability of the latest version, you can start to use the new functions available, which requires setting **OPMODE** to (NEWFUNC,800).

Include the procedures and applications you identified in steps 8 on page 487 and 9 on page 487 in your regression tests.

11. Review the tasks to customize z/OS, and the queue manager. Plan how to change the queue manager definitions and started task JCL to migrate your queue managers to the latest versions.

The customization steps for migration are described in “z/OS: Review and modify queue manager customizations from the previous release” on page 489

12. Review the usage of page set 0.

Issue the operator command **cpf, /cpf DISPLAY USAGE PSID(0)** to obtain a report on pageset 0 usage.

The size of queue definitions increased in IBM WebSphere MQ Version 7.1. During migration queue definitions stored on pageset 0 are rewritten if you are migrating from a previous release.

The rewrite is carried out as a single transaction when the queue manager is first migrated to IBM WebSphere MQ Version 7.1.

Ensure that there is enough space available on pageset 0 to create a copy of the queue definitions while migration is taking place. Typically, 60% free space on pageset 0 before migration is sufficient. However, the use of EXPAND(SYSTEM) on the pageset definition allows for automatic expansion as required.

If there is insufficient space on pageset 0 during migration, the queue manager abends with completion code X'5C6' and reason code X'00C91900'

13. Check that you are using a supported level of assembler or compiler.

You can write IBM MQ applications using any compiler capable of generating standard OS linkage to the IBM MQ stub routines.

Some of the data types used by IBM MQ API calls are not supported on some older compilers. You might require a more recent compiler. The following limitations are known:

- a. Assembler copy books contain blank lines, which are not tolerated by assemblers earlier than **HLASM**.
- b. Some older releases of PL/I do not support fixed bin(63) type. A macro defines such fields as char(8) when an earlier compiler is detected.
- c. Some older releases of COBOL do not support function-pointers, which are used by the MQCB API.

14. Plan any changes to libraries required by your applications and channel exits.

15. Plan to upgrade your IBM MQ MQI client installations to the latest version.

16. Plan to upgrade your client and server applications to use new functions in the latest version.

17. Plan to migrate other vendor software, such as WebSphere Application Server, or CICS to use the latest version.

Update the IBM MQ libraries in the STEPLIB and DFHRPL concatenations of your CICS region JCL and restart CICS.

**Note:**

- CICS

Update the IBM MQ libraries in the STEPLIB and DFHRPL concatenations of your CICS region JCL and restart CICS.



Up to, and including CICS 3.2, the connection between IBM MQ and CICS is provided by IBM MQ. You must change the SCSQCICS and SCSQAUTH libraries in the DFHRPL concatenation provided by IBM MQ.

After CICS 3.2, the connection between IBM MQ and CICS is provided by CICS libraries. Update the libraries, if you are using CICS Transaction Server for z/OS Version 3.2 or later. Without this change, you are not able to use the most recent IBM MQ features. You must change the SCSQCICS library in the DFHRPL concatenation provided by IBM MQ, and also the STEPLIB concatenation.

Create separate CICS started procedure JCL. For each CICS region that is connected to an IBM MQ queue manager, ensure that there is a separate CICS started procedure JCL.

This ensures that the modification of reference to a certain version of IBM MQ libraries in the CICS started procedure JCL only has impact for that single CICS region. In this way you can migrate one queue manager, and only the CICS region or regions connected to it, which makes staged migration possible.





CICS STEPLIB has thlqual.SCSQAUTH, and DFHRPL has thlqual.SCSQCICS, thlqual.SCSQLOAD, and thlqual.SCSQAUTH. For more information, see *Setting up the CICS- IBM MQ adapter*

18. Review any other installed SupportPacs for their applicability to the latest version.

## What to do next

Do the task, “z/OS: Review and modify queue manager customizations from the previous release” to migrate a queue manager. If you must restore a queue manager to the previous version, see “z/OS: Reverting a queue manager to a previous release” on page 563.

When you are confident existing applications are running without migration problems on the latest version, plan to update **OPMODE** to (NEWFUNC,800) to enable new function.

-  [CICS Transaction Server for z/OS, Version 3 Release 2 product documentation](#)
-  [The CICS-IBM MQ adapter](#)
-  [IBM MQ Support, Migration PTFs](#)
-  [IBM MQ - SupportPacs by Product](#)

## z/OS: Review and modify queue manager customizations from the previous release

Review the z/OS and IBM MQ customization steps, and change any customizations before starting any queue managers with the libraries of the latest version.

### Before you begin

You must change the JCL and configuration definitions for the queue manager when you migrate to the latest version. Make copies of the JCL libraries, configuration definitions, and the started task JCL. Modify the copied files to customize the queue manager definitions and JCL for the latest version.

1. Copy the JCL libraries containing the configuration definitions for your queue manager. You must change some of the members, for example to reference the libraries of the new version.
2. Copy the started task JCL for the queue manager and the channel initiator. You might change these members.

You can continue to run the queue manager on the previous version until it is ready to switch to the latest version. Preparing for the switch is a long process. Switching from the earlier version to the latest version is a quick process. The switch to the latest version occurs when you restart the queue manager:

- Switch to using your customized copies of the configuration definitions and started task JCL for the queue manager.

Create dataset aliases, such as MQM.MQP1.SCSLOAD, and reference them in JCL. Map the aliases to the real data sets, such as MQM.MQV71.SCSLOAD or MQM.MQV80.SCSLOAD.

Change the aliases to switch between the two sets of target libraries. With the aliases, you can start applications or the queue manager when moving to a new release of IBM MQ without changing the STEPLIB JCL; switch the aliases from referring to the earlier version of the product to refer to the latest version of the product.

- Until you have verified the start up, start the queue manager, channel initiator, and listener separately, checking for errors on the console after each component is started. If the start up runs cleanly, combine the start up of all three components in the production environment.

As part of this task you must set **OPMODE** to (COMPAT,800). If you have not set **OPMODE**, the queue manager does not start.

Note that if you use the supplied CSQ4ZPRM sample, the queue manager starts without any problems because the provided default value for **OPMODE** is COMPAT,800).

The (COMPAT,800) setting does two things:

1. It checks that you are starting the queue manager that links to the libraries for the latest version of the product.
2. It disables the use of new function in the latest version of the product.

You can revert to running the queue manager with the earlier version libraries, if you have never set **OPMODE** to (NEWFUNC,800)

**Tip:** Create data set aliases such as MQM.MQP1.SCSLOAD, and reference them in JCL. Map the aliases to the real data sets, such as MQM.MQV71.SCSLOAD or MQM.MQV80.SCSLOAD. Change the aliases to switch between the two sets of target libraries. With the aliases, you can start applications or the queue manager when moving to a new release of IBM MQ without changing the STEPLIB JCL.

**Tip:** You can use the z/OS command `D GrS,system,RES=(*,MQM.V701.SCSLOAD)` to display which jobs are using the specified data set, and so identify which jobs and JCL need to be changed.

## About this task

The steps in this task can be performed by restarting a queue manager once. You must restart the queue manager to apply essential maintenance, and to install early code before proceeding with migration.

The steps are based on the setup procedure for new queue managers; see Customizing your queue managers.

## Procedure

Steps 1 to 10 on page 492 are z/OS customization tasks.

Steps 11 on page 492 to 16 on page 493 are IBM MQ customization tasks.

1. Review that all the load libraries that must be APF authorized are authorized.

See Task 2: APF authorize the IBM MQ load libraries.

2. Update the LPA libraries with the new version of early code libraries.

See Task 3: Update the z/OS link list and LPA.

This step is unnecessary if you update the IBM MQ early code in all the LPARs as part of bringing your previous version IBM MQ libraries up to the latest maintenance level; see step 3 on page 491.

You must activate the early code with an IPL, or restart the queue manager after issuing the following command.

```
REFRESH QMGR TYPE(EARLY)
```

**Note:** As part of the update you are asked to stop the queue manager. While the queue manager is stopped, you must check the *qmgr*.REFRESH.QMGR security profile is set up, refresh the queue manager and restart it.

3. Make your configuration of the previous version ready for migration.

a. Apply current maintenance to the current version libraries.

Refer to the Preventive Service Planning (PSP) bucket for the current version of your product; see PSP Buckets - How to find them on Web.

b. Apply the latest version migration and toleration<sup>6</sup> PTFs to the previous version code; see IBM MQ Support, Migration PTFs.

c. If required by any of the PTFs, make Db2 configuration changes and rebind Db2.

See the appropriate section in Task 9: Select and set up your coupling facility offload storage environment.

d. Install the early code for the latest version.

Replace the early code.

e. Make the latest version early code and earlier version target libraries available on all the LPARs that are running queue managers.

If the queue manager is a member of a queue-sharing group, update all the systems in the group.

f. Restart IBM MQ systems.

The early code is activated by an IPL, or a by issuing the command REFRESH QMGR TYPE(EARLY), and restarting the queue manager.

g. Verify correct function before proceeding and review any special actions detailed in the PTFs.

If you require fall back at this stage, use normal maintenance procedures to revert to the code for the previous version before PTF application.

4. Update your procedure to start the queue manager.

Change the STEPLIB for the queue manager to reference the libraries of the new version of the product.

See Task 6: Create procedures for the IBM MQ queue manager.

IBM MQ uses z/OS memory objects above the bar for some functions. You must allow the queue manager to access storage above the bar.

Your installation might have customized the SMFPRMxx member of SYS1.PARMLIB, or the **IEFUSI** exit to provide a default limit for jobs using virtual storage above the 2 GB bar. Check these limits give sufficient memory for a queue manager. A reasonable starting allocation is 2GB. The message CSQY220I displays the amount of virtual storage currently used and available.

If your installation does not have a default limit for storage above the bar, or if you want to use a different limit for your queue manager, you can provide a queue manager-specific restriction on the amount of virtual storage available above the bar for memory objects by coding a **MEMLIMIT** parameter on the JCL of the queue manager stored procedure, xxxxMSTR, for example:

```
//PROCSTEP EXEC PGM=CSQYASCP,REGION=0M,MEMLIMIT=2G
```

MEMLIMIT defines memory available above the bar; see Address space storage for more information.

For specific information on setting MEMLIMIT, when using the accounting and statistics changes in IBM MQ Version 8.0, see Channel initiator storage usage.

5. Update your procedures for the channel initiator.

Change the STEPLIB for the channel initiator to reference the libraries of the new version of the product.

See Task 7: Create procedures for the channel initiator.

---

6. The “migration and toleration” PTFs are also known as the “backward migration and coexistence” PTFs. They are the same PTFs.

6. Review any automated alerts when queue manager and channel initiator messages, and errors, are detected. New messages might have been added that cause automated alerts, and some messages might have changed.
7. Review C language channel exits
 

Ensure your C language channel exits are using the following statement:

```
#pragma environment(function-name)
```

as defined in the C systems programming environment for system exits, described in the z/OS C/C++ Programming Guide.
8. Update the IBM MQ Db2 configuration.
 

If you are using queue-sharing groups, you must update the procedures.

Customize and run the CSQ4570T and CSQ4571T sample JCL in *hlq.SCSQPROC*.

Customize and run the CSQ45BPL and CSQ45GEX samples in *hlq.SCSQPROC*. Tailor these members to your environment, using your Db2 subsystem names and data set names.

CSQ45BPL of *hlq.SCSQPROC* contains the plan names required for the latest version of IBM MQ. CSQ45GEX of *hlq.SCSQPROC* contains the authorities required.

See steps 5 and 6 of Task 9: Select and set up your coupling facility offload storage environment.
9. Review your security controls for queue-sharing groups, the channel initiator, and all queue managers accessing the coupling facility list structures.
 

See Task 11: Implement your ESM security controls.

You must ensure that you have addressed the following points when you migrate your security profiles to the latest version of IBM MQ.

  - The External Security Manager software is at the correct version and level and that all of your prerequisite software is installed.
  - IBM MQ security classes have been updated to include the mixed case classes.
  - Enterprise has migrated to mixed case security; see “z/OS Migrating a queue manager to mixed case security” on page 561.
10. Change SYS1.PARMLIB to ensure that any changes you made dynamically remain in effect after an IPL.
 

SYS1.PARMLIB must reference the initialization input data sets that you customize in step 11.

See Task 12: Update SYS1.PARMLIB members.
11. Update the initialization input data sets.
 

Each IBM MQ queue manager gets its initial definitions from a series of commands contained in the IBM MQ initialization input data sets. These data sets are referenced by the Data Definition (DD) names CSQINP1 and CSQINP2 defined in the queue manager started task procedure.

See Task 13: Customize the initialization input data sets.

The CSQINP1 and CSQINP2 initialization input data sets include more samples and the contents of some samples have been moved to other samples. Particular changes to take note of are the commands to define queues to hold publish/subscribe state information. The commands must be in the right order.

You must review the customization you have made previously to CSQINP1 and CSQINP2, and merge them into the initial definitions provided with the latest version of the product.

Secure the server-connection channels used by clients; see Securing remote connectivity to the queue manager.

DEFINE SUB for SYSTEM.DEFAULT.SUB is no longer permitted in the CSQINP2 input data set. DEFINE SUB commands can instead be issued from the CSQINPT input data set. The CSQINPT input data set is processed each time the publish/subscribe engine is started, either during queue manager startup, or when the publish/subscribe engine is started with the **ALTER QMGR PSMODE(ENABLED)** command. See Issuing commands to IBM MQ for z/OS for more information on using the CSQINPT input data set.

12. Update your system parameter module.  
 Your system parameter module is based on the default module, CSQZPARM and CSQ4ZPRM.  
 See Task 17: Tailor your system parameter module.  
 Set the value of the **OPMODE** parameter in the **CSQ6SYSP** macro to (COMPAT, 800). The parameters (COMPAT, 800) force a check that the queue manager is started at the 800 command level, and is started in compatibility mode.
13. Update the libraries you added to STEPLIB concatenations to make Batch, TSO, and RRS adapters available to applications.  
 Change the STEPLIB for the Batch, TSO, and RRS adapters to reference the libraries of the new version of the product.  
 See Task 19: Set up Batch, TSO, and RRS adapters.

**Note:** You can connect applications that reference the latest version STEPLIB to a queue manager that is running on the latest version, or an earlier version. You must not connect applications that reference a STEPLIB from an earlier version to a queue manager running on a later version.

14. Update the libraries you added to connect CICS to the queue manager.  
 You must update the IBM MQ libraries in the STEPLIB and DFHRPL concatenations of your CICS region JCL and restart CICS. You are then able to use the most recent IBM MQ features.

**Note:**

- CICS

Update the IBM MQ libraries in the STEPLIB and DFHRPL concatenations of your CICS region JCL and restart CICS.

Up to, and including CICS 3.2, the connection between IBM MQ and CICS is provided by IBM MQ. You must change the SCSQCICS and SCSQAUTH libraries in the DFHRPL concatenation provided by IBM MQ.

After CICS 3.2, the connection between IBM MQ and CICS is provided by CICS libraries. Update the libraries, if you are using CICS Transaction Server for z/OS Version 3.2 or later. Without this change, you are not able to use the most recent IBM MQ features. You must change the SCSQCICS library in the DFHRPL concatenation provided by IBM MQ, and also the STEPLIB concatenation.

Create separate CICS started procedure JCL. For each CICS region that is connected to an IBM MQ queue manager, ensure that there is a separate CICS started procedure JCL.

This ensures that the modification of reference to a certain version of IBM MQ libraries in the CICS started procedure JCL only has impact for that single CICS region. In this way you can migrate one queue manager, and only the CICS region or regions connected to it, which makes staged migration possible.

CICS STEPLIB has thlqual.SCSQAUTH, and DFHRPL has thlqual.SCSQCICS, thlqual.SCSQLOAD, and thlqual.SCSQAUTH. For more information, see Setting up the CICS- IBM MQ adapter

15. Update the libraries you use to set up the operations and control panels.  
 Change the STEPLIB for the operations and control panel.  
 See Task 20: Set up the operations and control panels.

**Note:** You can connect the operations and control panel that references the latest version STEPLIB to the queue manager that is running on the latest version, or an earlier version. You must not connect the operations and control panel that references a STEPLIB from an earlier version to a queue manager running on a later version.

16. Update system libraries to format IBM MQ dumps using the Interactive Problem Control System (IPCS).  
 See Task 21: Include the WebSphere(r) MQ dump formatting member.

## What to do next

Return to the parent migration planning task, “z/OS: Migration planning to the latest release” on page 486.

### Related tasks:

“Migrating a single IBM MQ z/OS queue manager to the next release of the product” on page 558  
Carry out the instructions in this topic to migrate a single IBM MQ queue manager on z/OS,

## z/OS: IBM MQ Version 8.0 JCL changes

Tables showing new members and changed members between the IBM WebSphere MQ Version 7.1 hlq.SCSQPROC and the IBM MQ Version 8.0 hlq.SCSQPROC PDS libraries at general availability time.

### Notes:

1. Nearly all of the members do show changes between releases, but the majority of those changes are what are termed cosmetic. They are often simply changes to reflect the modification to the release number.
2. Any changes to the samples libraries since general availability are not included.
3. This information can be useful to your administrators when migrating existing queue managers, to make sure that new features are picked up correctly.

## New members

Table 72. IBM MQ for z/OS additional members

| Member name | Description                                                                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CSQ4AMSM    | Procedure to start the IBM MQ Advanced Message Security ( IBM MQ AMS ) task.                                                                                        |
| CSQ4BCNV    | Sample JCL to convert the BSDS to version 2 for the log RBA expansion.                                                                                              |
| CSQ4INSM    | CSQINP2 sample for SYSTEM objects for IBM MQ AMS.                                                                                                                   |
| CSQ4QLOD    | Sample JCL for the QLOAD utility.                                                                                                                                   |
| CSQ4SCHD    | Sample SCHEDxx member for the following resources: <ul style="list-style-type: none"><li>• queue manager</li><li>• channel initiator</li><li>• IBM MQ AMS</li></ul> |
| CSQ40CFG    | Sample JCL to define a new security policy for IBM MQ AMS. This member runs the setmqspl command                                                                    |
| CSQ40ENV    | Sample IBM MQ AMS environment variables.                                                                                                                            |
| CSQ40RSM    | Sample IBM MQ AMS audit report.                                                                                                                                     |
| CSQ45BPK    | Sample JCL to bind Db2 packages for shared queues. This job was added to IBM WebSphere MQ Version 7.1 using a program temporary fix.                                |
| CSQ4570T    | Sample job to migrate Db2 resource from IBM WebSphere MQ Version 7.0.0 and IBM WebSphere MQ Version 7.0.1 to IBM MQ Version 8.0.                                    |
| CSQ4571T    | Sample job to migrate Db2 resource from IBM WebSphere MQ Version 7.1 to IBM MQ Version 8.0.                                                                         |

## Modified members

Table 73. IBM MQ for z/OS changed members

| Member name              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CSQ4CHIN                 | STEPLIB concatenation changed to add the optional library<br>DSN=++CSFQUAL++.SCSFMOD0,DISP=SHR used when password protection is on.<br><br>For specific information on setting MEMLIMIT, when using the accounting and statistics changes in IBM MQ Version 8.0, see Channel initiator storage usage.                                                                                                                                                                            |
| CSQ4INPR and<br>CSQ4INP1 | Changed to support the new bufferpool definition attributes: <ul style="list-style-type: none"> <li>• LOCATION: Default is BELOW</li> <li>• PAGECLAS: Default is 4 KB</li> <li>• REPLACE: Default is NOREPLACE</li> </ul>                                                                                                                                                                                                                                                        |
| CSQ4INSG                 | Changed to add the: <ul style="list-style-type: none"> <li>• STATCHL attribute to the sample channel definitions</li> <li>• DEFINE AUTHINFO( 'SYSTEM.DEFAULT.AUTHINFO.IDPWOS' ) to support the new authorization feature</li> <li>• CLROUTE attribute to the sample default topic definitions</li> <li>• Definition for the SYSTEM.DDELAY.LOCAL.QUEUE for those using JMS 2.0</li> </ul>                                                                                         |
| CSQ4INSX                 | Added the model queue definition for the<br>SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE to support split cluster transmission queues.                                                                                                                                                                                                                                                                                                                                                    |
| CSQ4INYC                 | Added the STATCHL attribute to the sample cluster receiver.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CSQ4INYD                 | Trigger attributes on the sample transmission queue definition have been changed to: <pre>* Trigger attributes TRIGGER + TRIGTYPE( FIRST ) + TRIGMPRI( 0 ) + TRIGDPH( 1 ) + TRIGDATA( '++locqmgr++.TO.++remqmgr++' ) + PROCESS( ' ' ) + INITQ( 'SYSTEM.CHANNEL.INITQ' )</pre> <p>from</p> <pre>* Trigger attributes TRIGGER + TRIGTYPE( FIRST ) + TRIGMPRI( 0 ) + TRIGDPH( 1 ) + TRIGDATA( ' ' ) + PROCESS( '++remqmgr++.SEND.PROCESS' ) + INITQ( 'SYSTEM.CHANNEL.INITQ' )</pre> |
| CSQ4INYG                 | The sample ALTER QMGR statement has been modified to have all the attributes on an individual line.                                                                                                                                                                                                                                                                                                                                                                              |
| CSQ4INYR                 | A new storage class definition has been added: <pre>DEFINE STGCLASS( 'SYSLNGLV' ) + QSGDISP( QMGR ) + PSID( 02 )</pre>                                                                                                                                                                                                                                                                                                                                                           |
| CSQ4MSRR and<br>CSQ4MSTR | Several optional changes have been made: <ul style="list-style-type: none"> <li>• The MEMLIMIT parameter has been added to the procstep</li> <li>• The CSQINP2 concatenation now includes the new member CSQ4INSM (for IBM MQ AMS )</li> <li>• A new DD statement for CSQINPT has been added. This statement concatenates the members CSQ4INST and CSQ4INYT for publish/subscribe</li> <li>• A new output DD statement has been added, CSQOUTT</li> </ul>                        |

Table 73. IBM MQ for z/OS changed members (continued)

| Member name | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CSQ4PAGR    | The FORM EXEC statement has been altered to include a REGION=4M parameter. The CSQ4PAGE FORM EXEC statement already contained the REGION parameter.                                                                                                                                                                                                                                                                                                                                   |
| CSQ4SMJF    | New DD statements have been added in support of the new SMF data:<br><pre>//QCCT DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233) //QCTDSP DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233) //QCTADP DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233) //QCTSSL DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233) //QCTDNS DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233) //QCST DD SYSOUT=*,DCB=(LRECL=233,RECFM=F,BLKSIZE=233)</pre>                                                     |
| CSQ4ZPRM    | Several changes have been made: <ul style="list-style-type: none"> <li>• Archive log blocksize changed to 24576 from 28762</li> <li>• CONNSWAP=YES parameter added. This allows batch jobs to swap during an IBM MQ API call</li> <li>• EXCLMSG=( ) parameter added. This allows up to 16 messages to be suppressed.</li> <li>• OPMODE=(COMPAT,800) change from the IBM WebSphere MQ Version 7.1 values.</li> <li>• SPLCAP=NO, message encryption flag used by IBM MQ AMS.</li> </ul> |
| CSQ45BPL    | Sample job to bind Db2 plans using the Db2 TSO batch interface.<br><b>Note:</b> The plan names have changed for IBM MQ Version 8.0.                                                                                                                                                                                                                                                                                                                                                   |
| CSQ45CTB    | Columns added in the CSQ.ADMIN_B_STRBACKUP table:<br><pre>BSTART_RBA_HI CHAR(2) WITH DEFAULT X'0000', BEND_RBA_HI CHAR(2) WITH DEFAULT X'0000',</pre>                                                                                                                                                                                                                                                                                                                                 |
| CSQ456TB    | Changes to migrate the Db2 tables to IBM MQ Version 8.0: <ul style="list-style-type: none"> <li>• Added the two new columns (see the preceding entry for CSQ45CTB)</li> <li>• Table names are altered for IBM MQ Version 8.0.</li> </ul>                                                                                                                                                                                                                                              |
| CSQ45GEX    | Sample Db2 GRANT job to use the IBM MQ Version 8.0 names.                                                                                                                                                                                                                                                                                                                                                                                                                             |

## UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

### Before you begin

This scenario is one of three, which describe alternative ways to upgrade queue managers from an earlier version of the product. The other scenarios are as follows:

1. Install the latest version of the product alongside an earlier version ; see “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500.
2. Run the latest version of the product alongside an earlier version ; see “UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505.

Read these three tasks to plan how you are going to migrate to the multi-installation environment of the latest version. Even if you do not plan to have more than one version of the installation on a server, read this topic. The steps to upgrade IBM MQ from Version 7.0.1 to Version 8.0 have changed.



These topics are for planning multi-installation migration. The planning topics guide you in deciding what other tasks you must perform to migrate queue managers and applications to the latest version. For the precise sequence of commands to upgrade a queue manager to the latest version, do the migration task for the platform you are interested in. All the tasks are listed by platform in the links at the end of this topic. As part of the queue manager migration task, back up your existing queue manager data. Even on a multi-installation server, queue managers cannot be restored to a previous command level after migration.

## About this task

In the single-stage migration scenario, the installation of the latest version of the product replaces an earlier version in the same installation location. It is the same migration process that you would have used to upgrade the product prior to IBM WebSphere MQ Version 7.0.1.6 . It is now termed “single-stage” migration, in contrast to “side-by-side” and “multi-stage” migration.

The advantage of single-stage migration is that it changes the configuration of a queue manager on the earlier version as little as possible. Existing applications switch from loading the libraries from the earlier version, to loading the libraries of the latest version, automatically.

Queue managers are automatically associated with the installation on the latest version. Administrative scripts and procedures are affected as little as possible by setting the installation to be the primary installation. If you set the installation of the latest version to be the primary installation, commands such as **strmqm** work without providing an explicit path to the command.

Four types of object are important to consider during migration:

- Installations
- Queue managers
- Administrative procedures
- Applications


Administrative procedures contain IBM MQ commands, and scripts that use commands.

To run a command, the operating system must find the command in an IBM MQ installation. For some commands, you must run the command from the installation that is associated with the correct queue manager. IBM MQ does not switch commands to the correct installation. For other commands, such as **setmqinst**, you can run the command from any installation that has the latest version of the product installed.

If an earlier version of the product is installed, the command that is run is the command for that version, unless the search path is overridden by a local setting. You can override the search path by running **setmqenv**. If Version 7.0.1 is not installed, you must set the correct path to run a command. If you have set a primary installation, the command that is run is the copy in the primary installation, unless you override the selection with a local search path.


## Procedure


1. Stop local IBM MQ applications.
2. Stop all the queue managers and listeners. This scenario uses two queue managers QM1 and QM2.
3. Uninstall any fix packs you have installed from the previous IBM MQ version.
4. Upgrade the earlier version of the product to the latest version in the same installation directory.
  - a. Decide on an installation naming convention. Give the installation a name of your choosing, or accept the default installation name. For the first installation, the default name is *Installation1*. For the second installation, the name is *Installation2*, and so on.

 On AIX there is no option to set the installation name, *Installation1* is set by default.

- b. Upgrade the earlier version of the product to the latest version in place, or uninstall the earlier version, without deleting any queue managers, and install the latest version in the same default location. Whether you have to uninstall your previous version of the product depends upon your operating system.

On the following platforms, you do not have to uninstall a previous version of the product:

- AIX
- Windows
-  IBM i, where the process is known as a *slip* install

 If `mqm.xr.clients` and `mqm.txclient.rte` file sets from earlier versions are installed, you must uninstall these file sets from the earlier versions.

On the following platforms, you have to uninstall the previous version of the product:

- HP-UX
  - Linux
  - Solaris
- A reason for installing into the same location is to simplify application migration. If you change the installation location, you might remove IBM MQ libraries from an application search path. To migrate an application search path you must modify the application environment, or more rarely, the application itself.
  - The default installation path is specified as a load path in the IBM MQ build scripts for UNIX and Linux. After installation of the latest version, the load libraries of the latest version of IBM MQ are in the same location as were the libraries of the earlier version. If you built applications by following the examples in the product documentation for the earlier version, the applications load the correct libraries in the latest version.

5. Optional: Make the latest version of the installation the primary installation.

- a. Run the **setmqinst** command

On Windows

```
"Inst_1_INSTALLATION_PATH\bin\setmqinst" -i -n Inst_1
```

On UNIX

```
Inst_1_INSTALLATION_PATH/bin/setmqinst -i -n Inst_1
```

**Note:** Use the `dspmqinst` command to discover the <Installation name>, or use the default value <Installation 1>.

- Make the installation primary to avoid specifying a search path to run IBM MQ commands.
- If there is a primary installation, UNIX and Linux applications that expect to find the IBM MQ library in `/usr/lib`, find a symbolic link to the library in `/usr/lib/32`<sup>7</sup>. `/usr/lib/32` is normally in the default search path. It is also specified as a load path in the IBM MQ build scripts for UNIX and Linux.
- It is sufficient to link applications only to `/usr/lib`. With a primary installation of the latest version of the product defined on the server, an application can connect to any queue manager associated with any installation on the server. IBM MQ loads the correct library for the application.

6. Start the queue managers and applications.

- a. Optional: Run the **setmqm** command to associate the queue managers with `Inst_1`.

---

7. `/usr/lib` for 64 bit applications.

```
setmqm -m QM1 -n Inst_1
setmqm -m QM2 -n Inst_1
```

#### Notes:

- The **setmqm** step is optional only in the case where migration is from IBM WebSphere MQ Version 7.0.1 to a later release. In this case, the **strmqm** command automatically associates the queue manager with its own installation.
- If you are migrating between any other releases of IBM MQ, you must use **setmqm** to associate the queue managers with the new installation manually.

**Windows** If you have multiple installations, note that queue managers that were configured to start automatically, and remain after uninstalling IBM WebSphere MQ Version 7.0.1, automatically start under any other existing Version 7.1 (or later) installation when either the machine reboots, or the Service for that installation is restarted. In order to avoid this, ensure that all queue managers have been moved to the desired installation before uninstalling IBM WebSphere MQ Version 7.0.1.

- b. Run the **strmqm** command to start the queue managers and migrate them to the latest version of the product.

```
strmqm QM1
strmqm QM2
```

At this point, queue manager data is migrated and you cannot revert to a previous release.

- When an application connects to a queue manager, the operating system searches its load path to load the IBM MQ library<sup>8</sup>. A Version 7.1, or later, library contains code that checks that the queue manager is associated with an installation. If a queue manager is associated with a different installation, IBM MQ loads the correct IBM MQ library for the installation the queue manager is associated with.

## What to do next

You cannot reinstall an earlier version of the product on a system that has the latest, or any other, version of IBM MQ installed.

#### Related concepts:

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On **z/OS**, z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

#### Related tasks:

“IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485 Before migrating from one version to another, read this topic. Create your own migration plan based on the outline in the planning topic.

“UNIX systems - migrating a queue manager from your current version to the latest version” on page 513 Follow these instructions to migrate a queue manager from your current version to the latest version.

“Windows: Migrating a queue manager from a previous version to the latest version” on page 515 Follow these instructions to migrate a queue manager from your current version to the latest version.

---

8. On Windows, the IBM MQ library is a DLL. A DLL is sometimes called a load library or a shared library. The entry points to a DLL are defined in a link library, with the file extension .lib32 or .lib. The .lib library is linked at build-time and the DLL loaded at runtime.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

**Related information:**

Installing IBM MQ server on AIX

Installing IBM MQ server on HP-UX

Installing IBM MQ server on Linux

Installing IBM MQ server on Solaris

Installing IBM MQ server on Windows

Associating a queue manager with an installation

Changing the primary installation

Choosing an installation name

setmqenv

setmqinst

setmqm

---

## **UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version**

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

### **Before you begin**

If you are using IBM WebSphere MQ Version 7.0.1, you must ensure that you are running IBM WebSphere MQ Version 7.0.1.6 before installing the latest version of the product on the same server. Go to Fix Central to obtain the fix pack.

This scenario is one of three, which describe alternative ways to upgrade queue managers from an earlier version of the product. The other scenarios are as follows:

1. Replace the earlier version with the latest version; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496.
2. Run the latest version of the product alongside an earlier version ; see “UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505.

Read these three tasks to plan how you are going to migrate to the multi-installation environment of the latest version. The step-by-step migration scenario sits half-way between the single-stage and multi-stage migration scenarios.

These topics are for planning multi-installation migration. The planning topics guide you in deciding what other tasks you must perform to migrate queue managers and applications to the latest version. For the precise sequence of commands to upgrade a queue manager to the latest version, do the migration task for the platform you are interested in. All the tasks are listed by platform in the links at the end of this topic. As part of the queue manager migration task, back up your existing queue manager data. Even on a multi-installation server, queue managers cannot be restored to a previous command level after migration.

## About this task

In the “side-by-side” migration scenario, you install the latest version of IBM MQ alongside queue managers that continue to be associated with Version 7.0.1, or later.

When you are ready to migrate the queue managers, and applications, to the latest version:

1. Stop all the queue managers.
2. Uninstall the earlier version of the product.
3. Migrate all the queue managers and applications to the latest version.

The side-by-side migration scenario is less flexible than multi-stage migration, and might not seem to have any advantages over it. However, side-by-side migration does have advantages over the multi-stage and single-stage approaches. With the side-by-side approach, because you uninstall the earlier version before starting any queue managers, you can assign an installation on the latest version to be the primary installation.

In the multi-stage approach, you cannot set an installation of the latest version to be the primary installation while you continue to run the earlier version.

Having the latest version installation as the primary installation has two benefits.

1. With the latest version having the primary installation, many applications restart without reconfiguring their environment.
2. IBM MQ commands run against the primary installation, work without providing a local search path.

The advantage the side-by-side scenario has over the single-stage scenario is that you can install and verify the installation of the latest version of the product on the server before switching over to it.

Four types of object are important to consider during migration:

- Installations
- Queue managers
- Administrative procedures
- Applications


Administrative procedures contain IBM MQ commands, and scripts that use commands.

To run a command, the operating system must find the command in an IBM MQ installation. For some commands, you must run the command from the installation that is associated with the correct queue manager. IBM MQ does not switch commands to the correct installation. For other commands, such as **setmqinst**, you can run the command from any installation that has the latest version of the product installed.

If an earlier version of the product is installed, the command that is run is the command for that version, unless the search path is overridden by a local setting. You can override the search path by running **setmqenv**. If Version 7.0.1 is not installed, you must set the correct path to run a command. If you have set a primary installation, the command that is run is the copy in the primary installation, unless you override the selection with a local search path.

## Procedure

1. Install the latest version in a different installation directory from the earlier version.
  - a. Decide on an installation naming convention. Give the installation a name of your choosing, or accept the default installation name. For the first installation, the default name is *Installation1*. For the second installation, the name is *Installation2*, and so on.

 On AIX there is no option to set the installation name, *Installation1* is set by default.

- b. Verify the installation.

Run the installation verification procedures and your own tests.

2. Uninstall the earlier version of the product.

When uninstalling the earlier product, you must stop all queue managers and applications that have loaded an IBM MQ library on the server. For this reason, you might choose to postpone uninstalling the earlier version of the product until a convenient maintenance window. When an earlier version of the product is not installed on a server, it is sufficient to stop the queue managers and applications that have loaded libraries from the installation that you are uninstalling or updating. It is not necessary to stop applications and queue managers associated with other installations.

- a. Stop all applications that have loaded IBM MQ libraries on the server.
- b. Stop the queue managers and listeners on the server.
- c. Uninstall the earlier version of the product.
  - Stop all local IBM MQ applications
  - If you are migrating from IBM WebSphere MQ Version 7.0.1, stop all your queue managers and listeners, ensuring that you do not delete the queue managers.

**Note:** If you are migrating from a release other than IBM WebSphere MQ Version 7.0.1 you do not need to stop all the queue managers at this point.

3. Make the latest version of the installation the primary installation.

- a. Run the **setmqinst** command

On Windows

```
"Inst_1_INSTALLATION_PATH\bin\setmqinst" -i -n Inst_1
```

On UNIX

```
Inst_1_INSTALLATION_PATH/bin/setmqinst -i -n Inst_1
```

**Note:** Use the `dspmqinst` command to discover the <Installation name>, or use the default value <Installation 1>.

- Make the installation primary to avoid specifying a search path to run IBM MQ commands.

- If there is a primary installation, UNIX and Linux applications that expect to find the IBM MQ library in `/usr/lib`, find a symbolic link to the library in `/usr/lib/329`. `/usr/lib/32` is normally in the default search path. It is also specified as a load path in the IBM MQ build scripts for UNIX and Linux.
- It is sufficient to link applications only to `/usr/lib`. With a primary installation of the latest version of the product defined on the server, an application can connect to any queue manager associated with any installation on the server. IBM MQ loads the correct library for the application.

Use the `dspmqrinst` command to discover the `<Installation name>`, or use the default value `<Installation 1>`.

Doing this means that you do not have to specify a search path on IBM MQ commands.

#### 4. Start the queue managers and applications.

- Optional: Run the `setmqm` command to associate the queue managers with `Inst_1`.

```
setmqm -m QM1 -n Inst_1
setmqm -m QM2 -n Inst_1
```

#### Notes:

- The `setmqm` step is optional only in the case where migration is from IBM WebSphere MQ Version 7.0.1 to a later release. In this case, the `strmqm` command automatically associates the queue manager with its own installation.
- If you are migrating between any other releases of IBM MQ, you must use `setmqm` to associate the queue managers with the new installation manually.

**Windows** If you have multiple installations, note that queue managers that were configured to start automatically, and remain after uninstalling IBM WebSphere MQ Version 7.0.1, automatically start under any other existing Version 7.1 (or later) installation when either the machine reboots, or the Service for that installation is restarted. In order to avoid this, ensure that all queue managers have been moved to the desired installation before uninstalling IBM WebSphere MQ Version 7.0.1.

- Run the `strmqm` command to start the queue managers and migrate them to the latest version of the product.

```
strmqm QM1
strmqm QM2
```

At this point, queue manager data is migrated and you cannot revert to a previous release.

- When an application connects to a queue manager, the operating system searches its load path to load the IBM MQ library<sup>10</sup>. A Version 7.1, or later, library contains code that checks that the queue manager is associated with an installation. If a queue manager is associated with a different installation, IBM MQ loads the correct IBM MQ library for the installation the queue manager is associated with.

During this process you continue to use queue manager QM2 while you upgrade queue manager QM1 and you use queue manager QM1 while you upgrade QM2.

Note that each queue manager needs to be stopped in order to be associated with the new installation.

## What to do next

You cannot reinstall an earlier version of the product on a system that has the latest, or any other, version of IBM MQ installed.

9. `/usr/lib` for 64 bit applications.

10. On Windows, the IBM MQ library is a DLL. A DLL is sometimes called a load library or a shared library. The entry points to a DLL are defined in a link library, with the file extension `.lib32` or `.lib`. The `.lib` library is linked at build-time and the DLL loaded at runtime.

“IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485

Before migrating from one version to another, read this topic. Create your own migration plan based on the outline in the planning topic.

Installing IBM MQ server on AIX

Uninstalling IBM MQ on AIX

Uninstalling IBM MQ on HP-UX

Installing IBM MQ server on HP-UX

Uninstalling IBM MQ on Linux

Installing IBM MQ server on Linux

Uninstalling IBM MQ on Solaris

Installing IBM MQ server on Solaris

Uninstalling IBM MQ on Windows systems

Installing IBM MQ server on Windows

Associating a queue manager with an installation

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.


“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

Changing the primary installation

Choosing an installation name

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

setmqenv



setmqinst  
setmqm

---

## UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

### Before you begin

If you are using IBM WebSphere MQ Version 7.0.1, you must ensure that you are running IBM WebSphere MQ Version 7.0.1.6 before installing the latest version of the product on the same server. Go to Fix Central to obtain the fix pack.

This scenario is one of three, which describe alternative ways to upgrade queue managers from an earlier version of the product. The other scenarios are as follows:

1. Replace the earlier version with the latest version; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496.
2. Install the latest version of the product alongside an earlier version ; see “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500.

Read these three tasks to plan how you are going to migrate to the multi-installation environment of the latest version. The multi-stage migration scenario is the most flexible approach to migrating from IBM WebSphere MQ Version 7.0.1 to the latest version.

These topics are for planning multi-installation migration. The planning topics guide you in deciding what other tasks you must perform to migrate queue managers and applications to the latest version. For the precise sequence of commands to upgrade a queue manager to the latest version, do the migration task for the platform you are interested in. All the tasks are listed by platform in the links at the end of this topic. As part of the queue manager migration task, back up your existing queue manager data. Even on a multi-installation server, queue managers cannot be restored to a previous command level after migration.

#### Note:

- If an application uses COM or ActiveX it can connect to any queue manager as long as there is a primary installation and it is Version 7.1 or later.
- If you are running the IBM MQ.NET monitor in transactional mode, the queue manager it connects to must be the primary installation.

You cannot migrate these applications to the latest version until you uninstall the earlier version.

### About this task

In the multi-stage migration scenario, you install the latest version of the product alongside running queue managers that continue to be associated with the earlier version. You can create queue managers and run new applications using the latest version installation. When you are ready to start migrating queue managers and applications from the earlier, you can do so, one-by-one. When migration to the latest version is complete, uninstall the earlier version, and make the latest version installation the primary installation.

With the multi-stage approach, until you uninstall the earlier version, you must configure an environment to run applications that connect to a queue manager to the latest version. You must also provide a path to run IBM MQ commands. Both these tasks are accomplished with the **setmqenv** command.

**Note:** When you have uninstalled the earlier version, and set the latest version as a primary installation, in most circumstances it is not necessary to run the **setmqenv** command to run applications. It is still necessary to run **setmqenv** to set the environment for commands that connect to a queue manager associated with an installation that is not primary.

Four types of object are important to consider during migration:

- Installations
- Queue managers
- Administrative procedures
- Applications


Administrative procedures contain IBM MQ commands, and scripts that use commands.

To run a command, the operating system must find the command in an IBM MQ installation. For some commands, you must run the command from the installation that is associated with the correct queue manager. IBM MQ does not switch commands to the correct installation. For other commands, such as **setmqinst**, you can run the command from any installation that has the latest version of the product installed.

If an earlier version of the product is installed, the command that is run is the command for that version, unless the search path is overridden by a local setting. You can override the search path by running **setmqenv**. If Version 7.0.1 is not installed, you must set the correct path to run a command. If you have set a primary installation, the command that is run is the copy in the primary installation, unless you override the selection with a local search path.

## Procedure

1. Install the latest version in a different installation directory from the earlier version and verify the installation.
  - a. Decide on an installation naming convention. Give the installation a name of your choosing, or accept the default installation name. For the first installation, the default name is *Installation1*. For the second installation, the name is *Installation2*, and so on.

 On AIX there is no option to set the installation name, *Installation1* is set by default.

- b. Verify the installation.

Run the installation verification procedures and your own tests.

- You might create new queue managers running the latest version, and start to develop new applications before migrating applications from the earlier version.

2. Configure the operating system so that applications load the libraries for the latest version of the product.

- a. Migrate queue managers one at a time.

The first set of applications to load the libraries for the latest version of the product are the applications that connect to the first queue manager you are going to migrate.

It does not matter if those applications also connect to other queue managers on the server. If the applications load the latest version libraries, IBM MQ automatically loads the libraries for the earlier version for those applications that connect to that version.

You can either migrate the operating system environment of all applications, or just the applications that connect to the first queue manager you are going to migrate.

- b. Migrate IBM MQ MQI client applications

Some of the applications might be running as IBM MQ MQI client applications on another workstation. When you migrate a queue manager, clients connected to it continue to run without loading a client library for the latest version.

You can migrate these clients later, when you need to do so.

**Important:** If any IBM MQ MQI client applications are using the library for the earlier version on the server, you must eventually migrate the clients to use the latest version of the product before you uninstall the earlier version.

3. Migrate an application to load the new library for the latest version:

- Run **setmqenv** to modify the local path that is searched for IBM MQ libraries.
- Modify the global search path that is searched for IBM MQ libraries.
- Relink applications with an additional runtime load path.

Consult operating system documentation about how to modify the global search path, or include a fixed runtime load path in the application load module.

To run **setmqenv** using the **-s** option:

Windows:

```
"Inst_1_INSTALLATION_PATH\bin\setmqenv" -s
```

The **-s** option sets up the environment for the installation that runs the **setmqenv** command.

UNIX:

```
.Inst_1_INSTALLATION_PATH/bin/setmqenv -s -k
```

The **-k** option inserts the path to the IBM MQ load libraries at the start of the **LD\_LIBRARY\_PATH** environment variable, and adds the variable to the local environment; see "Loading IBM MQ libraries" on page 461.

**Note:** On UNIX the leading "." is critical. The dot followed by a space instructs the command shell run **setmqenv** in the same command shell and inherit the environment set by **setmqenv**.

4. Restart the queue manager and the applications that connect to it.
- a. Set up the local environment to the installation **Inst\_1**.

Windows:

```
"Inst_1_INSTALLATION_PATH\bin\setmqenv" -s
```

The **-s** option sets up the environment for the installation that runs the **setmqenv** command.

UNIX:

```
.Inst_1_INSTALLATION_PATH/bin/setmqenv -s
```

- b. Run the **setmqm** command to associate QM1 with **Inst\_1**.

```
setmqm -m QM1 -n Inst_1
```

- c. Run the **strmqm** command to start QM1 and migrate it to the latest version.

strmqm QM1

d. Restart application 1

The application loads the latest version library and connects to QM1, which is associated with the latest version of the product.

5. Migrate all queue managers and applications to the latest version.

Repeat steps 2 on page 506 and 4 on page 507, when required, until all the queue managers and applications are migrated to the latest version of the product.

6. Uninstall the earlier version of the product.

When uninstalling the earlier product, you must stop all queue managers and applications that have loaded an IBM MQ library on the server. For this reason, you might choose to postpone uninstalling the earlier version of the product until a convenient maintenance window. When an earlier version of the product is not installed on a server, it is sufficient to stop the queue managers and applications that have loaded libraries from the installation that you are uninstalling or updating. It is not necessary to stop applications and queue managers associated with other installations.

a. Stop all applications that have loaded IBM MQ libraries on the server.

b. Stop the queue managers and listeners on the server.

c. Uninstall the earlier version of the product.

- Stop all local IBM MQ applications
- If you are migrating from IBM WebSphere MQ Version 7.0.1, stop all your queue managers and listeners, ensuring that you do not delete the queue managers.

**Note:** If you are migrating from a release other than IBM WebSphere MQ Version 7.0.1 you do not need to stop all the queue managers at this point.

7. Make Inst\_1 the primary installation.

a. Run the **setmqinst** command

On Windows

```
"Inst_1_INSTALLATION_PATH\bin\setmqinst" -i -n Inst_1
```

On UNIX

```
Inst_1_INSTALLATION_PATH/bin/setmqinst -i -n Inst_1
```

**Note:** Use the `dspmqinst` command to discover the <Installation name>, or use the default value <Installation 1>.

- You do not have to set up a search path to run IBM MQ commands from the primary installation.
- If you set an installation of the latest version of the product as primary on UNIX and Linux, you do not have to set up `LD_LIBRARY_PATH` in most cases. You can remove calls to **setmqenv** to set `LD_LIBRARY_PATH`.


## What to do next

You cannot reinstall an earlier version of the product on a system that has the latest, or any other, version of IBM MQ installed.

Now that you have uninstalled the earlier version of the product, and made the latest installation primary, you can review how the application runtime environment is set. It is no longer necessary to run **setmqenv** to set up the search path to load libraries for the latest version. If you have only one installation of the latest version of the product installed, it is not necessary to run **setmqenv** to run commands.

**Related concepts:**

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

**Related tasks:**

“IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485  
Before migrating from one version to another, read this topic. Create your own migration plan based on the outline in the planning topic.

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

**Related information:**

Installing IBM MQ server on AIX

Installing IBM MQ server on HP-UX

Installing IBM MQ server on Linux

Installing IBM MQ server on Solaris

Installing IBM MQ server on Windows

Associating a queue manager with an installation

Changing the primary installation

Choosing an installation name

setmqenv

setmqinst

setmqm

---

## Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0

Migrate IBM MQ Telemetry from Version 7.0.1 to Version 8.0 by completing the tasks in this section. You must stop all IBM MQ activity on the system before migrating.

### About this task

In IBM WebSphere MQ Version 7.0.1, IBM MQ Telemetry was a separate feature. Because IBM MQ Telemetry is a component of IBM WebSphere MQ Version 7.1 and later, upgrading is essentially uninstalling IBM MQ Telemetry and installing a later version of IBM MQ. IBM MQ Telemetry can either be installed with the main product, or installed after the main product has been installed.

After the successful upgrade, Linux systems retain all telemetry data kept in `/var/mqm`, and the Windows systems retain the telemetry data in the installation directory of the product, for example: `C:\Program Files (x86)\WebSphere MQ`. Telemetry data is migrated to the later version of the product when the queue manager is started again.

#### Notes:

1. You can administer IBM MQ Telemetry Version 7.0.1 only from the Version 7.0.1 MQ Explorer. If you connect the Version 8.0 explorer remotely to a Version 7.0.1 queue manager, no telemetry resources are displayed. You cannot connect a Version 8.0 MQ Explorer locally to a Version 7.0.1 queue manager on the same server.
2. From IBM MQ Version 8.0, the Client Software Development Kit (SDK) is no longer supplied as part of the product. Instead, the current version of the SDK is available as the free download IBM Messaging Telemetry Clients SupportPac.
3. If you have upgraded from a previous version of the product you must download the latest version of the SDK.

The migration tasks, depending on your platform, are described in the subtopics.

#### Related information:

Installing IBM MQ

Installing IBM MQ Telemetry

## Windows: Migrating IBM MQ Telemetry from Version 7.0.1 to Version 8.0

Follow these instructions to migrate IBM MQ Telemetry from Version 7.0.1 to Version 8.0 on Windows.

### Before you begin

Before proceeding with this task, ensure that you back up your existing IBM MQ installation. You must stop the IBM MQ Telemetry service `SYSTEM.MQXR.SERVICE` before migrating.

### About this task

This task outlines the steps necessary to migrate your existing installation of IBM MQ Telemetry to Version 7.5 on Windows systems.

### Procedure

1. Uninstall IBM WebSphere MQ Version 7.0.1 using the control panel, following this procedure:
  - a. From the Windows task bar, click **Start > Control Panel**. (On some systems, click **Start > Settings > Control Panel** )
  - b. When uninstalling IBM MQ Telemetry by using the control panel, on Windows 7 systems, click **Programs and Features**, on other Windows systems, click **Add or Remove Programs**.

- c. Click IBM MQ Telemetry, then click **Change/Remove**. The uninstaller starts and summarizes items to be uninstalled.
  - d. Click **Uninstall**. The uninstaller lists the items being uninstalled. The Uninstall Complete screen displays the status of the uninstall. Click **Done**
2. Verify that the data folders still exist. For example, they might be located at C:\Program Files (x86)\WebSphere MQ.
  3. IBM MQ Telemetry is installed in one of two ways:
    - Installed as part of a **Custom** installation at the same time as IBM WebSphere MQ Version 7.0.1 during migration: Begin here: “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485
    - Added at a later date to an existing installation of IBM WebSphere MQ Version 7.5 or later: Begin here: Installing IBM MQ Telemetry
  4. Verify that the IBM MQ Telemetry migration was successful. See Verifying the installation of IBM MQ Telemetry .

## Results

Message AMQ4616 indicates completion of the task. The existing MQTT channels and previous subscriptions are still present.

### Related information:

Installing IBM MQ Telemetry

Verifying the installation of IBM MQ Telemetry

Verifying the installation of IBM MQ Telemetry by using IBM MQ Explorer

## Linux: Migrating from IBM MQ Telemetry Version 7.0.1 to Version 8.0

Follow these instructions to migrate IBM MQ Telemetry from Version 7.0.1 to Version 8.0 on Linux.

### Before you begin

Before proceeding with this task, ensure that you back up your existing IBM MQ installation. You must stop the IBM MQ Telemetry service `SYSTEM.MQXR.SERVICE` before migrating.

### About this task

This task outlines the steps necessary to migrate your existing installation of IBM MQ Telemetry to version Version 7.5 on Linux systems.

### Procedure

1. Uninstall IBM MQ Telemetry Version 7.0.1, following this procedure:
  - a. Navigate to the uninstallation directory. The default location on Linux is `/opt/mqm/mqxr/Uninstall_MQTT`.
  - b. Start the uninstaller, using the executable or binary file. On Linux systems, run `./Uninstall_MQTelemetry -i GUI`. The uninstaller starts and summarizes what is to be uninstalled.
  - c. Click **Uninstall**. The uninstaller lists the items being uninstalled.
  - d. Click **Done**.
2. Verify that the data folders still exist. Consult your `MQ_INSTALLATION_PATH` environment variable to find these folders.
3. IBM MQ Telemetry is installed in one of two ways:
  - Installed as part of a **Custom** installation at the same time as IBM WebSphere MQ Version 7.0.1 during migration: Begin here: “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485

- Added at a later date to an existing installation of IBM WebSphere MQ Version 7.5 or later: Begin here: Installing IBM MQ Telemetry
4. Verify that the IBM MQ Telemetry migration was successful. See Verifying the installation of IBM MQ Telemetry .

## Results

Message AMQ4616 indicates completion of the task. The existing MQTT channels and previous subscriptions are still present.

### Related information:

Installing IBM MQ Telemetry

Verifying the installation of IBM MQ Telemetry

Verifying the installation of IBM MQ Telemetry by using IBM MQ Explorer

---

## Migrating a queue manager to the latest release

The procedures for migrating a queue manager to the latest release are detailed in the following topics.

### Before you begin

This information applies to migrating a queue manager on platforms other than z/OS.


If you have installed early support program code on the server, you must delete all the queue managers created with the installation. Uninstall the code before proceeding with installing the production level code.

### About this task

Migration is a complex task. It goes beyond upgrading the IBM MQ code. The term “upgrading” applies to the process of installing a new code level. “Migrating” refers to the process of upgrading code and the to the task of migrating queue manager data, such as messages, queues, channels, and other resources. Migrating queue manager data is automatic.

### Procedure

Queue manager migration follows this overall plan:

1. Review the IBM MQ system requirements for the latest version; see IBM MQ System Requirements.
2. Back up your system. In particular back up the queue manager.  
You cannot run a queue manager on a previous version of IBM MQ, once you start the queue manager on the new version.
3. Carry out the platform-specific migration or upgrade procedure.  
Consult “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485.
  - a. Upgrade IBM MQ.
  - b. Customize the configuration.  z/OS  
Configuration changes are typically required only on z/OS.
  - c. Verify the installation.  
Create a queue manager to verify the installation. Verify applications and management tasks work with the new level of IBM MQ before migrating existing queue managers.
4. Perform any additional migration tasks that are required.  
If you are using publish/subscribe, you must migrate the publish/subscribe broker.



If the queue manager is a member of a queue manager cluster, or queue-sharing group, migrate the other members of the cluster or group.

**Important:** You must migrate the publish/subscribe broker state before you migrate your IBM MQ system to IBM MQ Version 8.0, as broker publish/subscribe migration is not supported in IBM MQ Version 8.0.

 [IBM MQ System Requirements](#)

## UNIX systems - migrating a queue manager from your current version to the latest version

Follow these instructions to migrate a queue manager from your current version to the latest version.

### Before you begin

1. The upgrade from your current version to the latest version of the product requires a full migration of queue managers. Create a migration plan. Use the planning task, [Planning migration to the latest version](#), as a guide.

2. Review the IBM MQ system requirements for the latest version; see [IBM MQ System Requirements](#).

3. Back up your system before you install the latest release of IBM MQ over a previous release. Once you have started a queue manager you cannot revert to the previous release.

If you must restore the system, you cannot recover any work, such as changes to messages and objects, performed by the latest version of IBM MQ. For more information about backing up your system, see [Backing up and restoring IBM MQ queue manager data](#).

4. Review any other installed SupportPacs for their applicability to the latest release.

5. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see [setmqenv](#).

### About this task

Complete this task to migrate a queue manager to the latest version of IBM MQ from an earlier version.

### Procedure

1. Log in as a user in group mqm.

2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their SYSTEM.FTE.STATE queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.

a. Run the **dspmqs** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqs -o installation -o status
dspmqs -a
```

**dspmqs -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqs -a** displays the status of active queue managers associated with the installation from which the command is run.

b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm lsr -m QMgrName
```

4. Back up the queue manager.
5. Log in as root.
6. If you are running IBM WebSphere MQ Version 7.0.1.6 or later, optionally uninstall the current version of IBM MQ. Note, that you carry out this step only if you are doing a single stage migration; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496
7. If you are running IBM WebSphere MQ Version 7.0.1.5 or earlier, uninstall the current version of IBM MQ. If you require a side-by-side migration or multistage migration only, you must upgrade to IBM WebSphere MQ Version 7.0.1.6
8. Install the latest version of the product. See the appropriate topic for the platform that your enterprise uses:
  - Installing IBM MQ Server on AIX .
  - Installing IBM MQ Server on HP-UX.
  - Installing IBM MQ Server on Linux .
  - Installing IBM MQ Server on Solaris.
9. Move the queue manager to the new IBM MQ installation. You need to carry out this step, only if you are running IBM WebSphere MQ Version 7.0.1.6 or later, and did not uninstall your current version of IBM MQ.
 

See “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500 or “UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505 for further information.
10. Start the queue manager.

```
strmqm QmgrName
```

When you first start a queue manager after migration:

- Any new attributes for existing objects are set to their default values.
- Any new default objects are created.

- Queue manager data is migrated.

**Important:** Do not use the -c option to start the queue manager, unless you explicitly want to reset or recreate the default system objects.

You must start IBM MQ before you start any listeners.

Backing up and restoring a queue manager

“The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423  
IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

## Windows: Migrating a queue manager from a previous version to the latest version

Follow these instructions to migrate a queue manager from your current version to the latest version.

### Before you begin

1. Create a migration plan; see Planning migration to the latest version
2. For more information about the versions of Windows that IBM MQ Version 8.0 supports, see System Requirements for IBM MQ V8.0.
3. Review the IBM MQ system requirements for the latest version; see IBM MQ System Requirements.
4. Back up your system before you install the latest release of IBM MQ over a previous release. Once you have started a queue manager you cannot revert to the previous release.

If you must restore the system, you cannot recover any work, such as changes to messages and objects, performed by the latest version of IBM MQ. For more information about backing up your system, see Backing up and restoring IBM MQ queue manager data.

5. Review any other installed SupportPacs for their applicability to the latest release.
6. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see setmqenv.

### About this task

Complete this task to migrate a queue manager to the latest version of IBM MQ from an earlier version.

All the objects that you previously created are maintained. The components that were previously installed are preselected in the feature options when you install the new level. If you leave these components selected, you can keep them or reinstall them. If you clear any of these components, the installation process uninstalls them. By default, a typical migration installs only the same features that were installed in the previous version installation.

For example, if MQ Explorer was not installed in an earlier installation, it is not stored in a later installation. If you want MQ Explorer, select a custom installation, and select the MQ Explorer feature on the Features panel. If you do not want MQ Explorer, uninstall the MQ Explorer feature by selecting a custom installation. Then clear the MQ Explorer feature on the Features panel. For more information about how to uninstall features, see Modifying the installation using IBM MQ Installation Launchpad.

You can also migrate a queue manager to a later release, on a system where a previous version has been uninstalled. The queue manager data must have been retained, or restored from a backup.

### Procedure

1. Log in as a user in group mqm.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their SYSTEM.FTE.STATE queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the **dspmqr** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqr -o installation -o status
dspmqr -a
```

**dspmqr -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm|sr -m QMgrName
```

4. Back up the queue manager.
5. Stop the MQSeries Service and exit the Service icon application.
6. Modify the IBM MQ installation, if necessary, by using one of these procedures:
  - Modifying the installation using IBM MQ Installation Launchpad
  - Silently modifying an IBM MQ server installation using **msiexec**
7. Reenter domain, user ID, and password information

When the installation of the latest version completes, the Prepare WebSphere MQ Wizard starts automatically.

**Where UAC is enabled:** If you rerun the Prepare WebSphere MQ Wizard, ensure that the wizard is run with Administrator privilege, otherwise the wizard might fail.

8. If you are running IBM WebSphere MQ Version 7.0.1.6 or later, optionally uninstall the current version of IBM MQ. Note, that you carry out this step only if you are doing a single stage migration; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

## What to do next

You might be restoring a previous version on a server with multiple IBM MQ installations. If one of the installations is primary, after restoring the previous version that installation, by default, becomes the primary installation.

You must review how applications connect to an installation. After restoring the previous version, some applications might connect to the wrong installation.

 [Fix Central](#)

 [IBM Passport Advantage](#)

“Migrating a queue manager in a high-availability configuration” on page 481

Follow standard procedures to migrate a queue manager that is part of a high-availability configuration. (On platforms other than z/OS.)

“Queue manager cluster migration” on page 478

You can migrate queue managers in a cluster all at once, or one at a time, which is called a staged migration. Migrate full repository queue managers in a cluster before partial repository queue managers.

“Reverting a queue manager to a previous version” on page 473

You can remove an upgrade before you have started a queue manager. After a queue manager has been started, if you remove the upgrade, the queue manager will not work.

 [IBM MQ - SupportPacs by Product](#)

“Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )” on page 453

You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

“IBM MQ migration” on page 448

Migration is the conversion of programs and data to work with a new code level of IBM MQ. Some types of migration are required, and some are optional. Queue manager migration is never required after applying a maintenance level update, that does not change the command level. Some types of migration are automatic, and some are manual. Queue manager migration is typically automatic and required after releases and manual and optional after a maintenance level upgrade that introduces a new function. Application migration is typically manual and optional.

“IBM MQ upgrades and fixes” on page 442

The term upgrade applies to changing the version V, release R, or modification M of a product. The term fix applies to a change in the F digit.

## IBM i Migrating a queue manager from a previous release



Follow these instructions to migrate a queue manager from a previous release to the latest release.

### Before you begin

If you decide to do a side-by-side installation, you must prepare the new server first, installing the prerequisite software.

1. Create a migration plan. Use the planning task, Planning migration to the latest version, as a guide.
2. Review the IBM MQ system requirements for the latest release of the product; see IBM MQ System Requirements
3. Review any other installed SupportPacs for their applicability to the latest release of IBM MQ.

## About this task

There are various types of migration:

- The migration takes place on the same machine, optionally accompanied by a hardware upgrade. This migration is referred to as a *slip install*.
- The migration takes place on a different machine. This migration is referred to as a *side-by-side install*.

A side-by-side installation gives you the option of preparing the new environment first, without interrupting the queue manager. It also gives you the limited option of reverting to use the previous release installation, if the migration is unsuccessful. It is limited, because you cannot restore the queue manager data from the latest version. You must restart processing with the queue manager data at the point you stopped the queue manager on the previous release.

If you want to add IBM MQ Advanced Message Security to your system, you must select Option (2) when you install the product; see Installing IBM MQ Advanced Message Security on IBM i for further information.

## Installation methods on IBM i



Select a slip installation or a side-by-side installation to upgrade IBM MQ for IBM i.

### About this task

A slip installation upgrades IBM MQ for IBM i on a computer with an earlier version is installed.

A side-by-side installation upgrades IBM MQ for IBM i on a different computer. You must save your queue managers before you start.

Follow the steps in the following tasks to carry out an upgrade.

The steps for both forms of upgrade are identical, except that you do not carry out the actions described in "Restore queue managers after upgrading IBM MQ on IBM i" on page 526 for a slip install.

## End IBM MQ activity on IBM i



End IBM MQ applications and connections, and remove any unwanted or indoubt messages.

### About this task

Before performing a slip install or side-by-side install, carry out the following procedure:

#### Procedure

1. Sign on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.
2. Stop all applications that are using the existing version of IBM MQ. Use the command WRKMQM, option 22, Work with queue manager jobs, to help find them.
3. End all channels for all queue managers on the system. To do so, use the WRKMQMCHL command and select option 15.

4. On each queue manager, end the command server. To do so, enter the command:  
`ENDMQMCSVR MQMNAME( QMGRNAME ) OPTION(*IMMED)`  
 where *QMGRNAME* is the name of the queue manager.
5. Remove any unwanted messages from your queues.
6. Resolve any in-doubt messages that are held by sender or server channels. To do so, use the `WRKMQMCHST` command and select option 17.
7. On each queue manager, save the latest media recovery checkpoint. To do so, enter the following command:  
`RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME( QMGRNAME ) DSPJRNDTA(*YES)`

## Quiesce IBM MQ on IBM i



Stop all queue managers. If necessary force all queue managers to stop, tidy up shared memory and end all jobs in the QMQM subsystem.

### About this task

The orderly shutdown of IBM MQ is called *quiescing*. You need to quiesce IBM MQ to upgrade to a newer version.

To quiesce one or more queue managers:

### Procedure

1. Sign on to a new interactive IBM i session, ensuring that you are not accessing any IBM MQ objects.
2. Ensure that you have:
  - a. \*ALLOBJ authority, or object management authority for the QMQM library.
  - b. Sufficient authority to use the ENDSBS command.
3. Warn all users that you are going to stop IBM MQ.
4. Use the `ENDMQM` command to quiesce all queue managers:  
`ENDMQM MQMNAME(*ALL) OPTION(*CNTRLD) ENDCCTJOB(*YES) RCDMQMIMG(*YES)  
 TIMEOUT( 15 )`

Where *15* is a timeout value in seconds.

If the `ENDMQM` command has not completed within a reasonable period (at least 10 minutes), use the `WRKMQM` command. This command identifies the queue managers that are still ending. Then force each one in turn to stop by issuing:

```
ENDMQM MQMNAME(QMGRNAME) OPTION(*IMMED)
```

Where *QMGRNAME* is the name of the queue manager.

Complete the tidying up of shared memory by issuing the command:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED) ENDCCTJOB(*YES) RCDMQMIMG(*NO)

TIMEOUT(15)
```

5. If the command in the previous step does not complete, end the subsystem immediately by issuing:  
`ENDSBS SBS(QMQM) OPTION(*IMMED)`
6. If this command also fails, use the operating system command `ENDJOB` to end all jobs in the subsystem QMQM, as described in the following steps.

**Note:** Do not use `ENDJOBABN` unless you intend to perform an IPL on the machine before starting IBM MQ. Ending IBM MQ jobs using `ENDJOBABN` can lead to damaged semaphores, which in turn can prevent your queue manager from starting.

a. If a QMGR must be shut down manually, the recommended order of ending jobs (ENDJOB) is shown in the list that follows. Wait a few minutes for AMQA\* or AMQZ\* jobs to tidy up.

- 1) RUNMQLSR - TCP listener (multi-threaded)
- 2) AMQCLMAA - TCP listener (single-threaded)
- 3) AMQRMPPA - Channel process pooling job
- 4) RUNMQCHI - channel initiator
- 5) AMQCRSTA - receiving MCA jobs
- 6) RUNMQCHL - sending MCA jobs
- 7) AMQCRS6B - LU62 receiver channel
- 8) AMQPCSEA - command server
- 9) RUNMQTRM - Application trigger monitor
- 10) RUNMQDLQ - Dead letter queue handler
- 11) AMQFCXBA - IBM Integration Bus Worker Job
- 12) AMQFQPUB - Queued Publish/Subscribe Daemon
- 13) RUNMQBRK - IBM Integration Bus Control Job
- 14) AMQZMUC0 ('0' is a zero) - Utility Manager
- 15) AMQZMUF0 ('0' is a zero) - Utility Manager
- 16) AMQZMUR0 ('0' is a zero) - Utility Manager
- 17) AMQZMGR0 ('0' is a zero) - Process Controller
- 18) AMQRRMFA - cluster repository manager
- 19) AMQZDMAA - deferred message manager
- 20) AMQALMPX - Log Manager
- 21) AMQZFUMA - object authority manager
- 22) AMQZLSA0 ('0' is a zero) - LQM agents
- 23) AMQZLAA0 ('0' is a zero) - LQM agents
- 24) AMQZXMA0 ('0' is a zero) - Execution Controller

b. Issue the following command:

```
ENDMQM MQMNAME(QMGRNAME) OPTION(*IMMED)
```

c. Issue the following command:

```
ENDMQM MQMNAME(*ALL) OPTION(*CNTRLD) ENDCCTJOB(*YES) RCDMQMIMG(*NO)
TIMEOUT(05)
```

Where 05 is a timeout value in seconds.

d. Manually clean up shared memory. Issue the following command:

```
EDTF '/QIBM/UserData/mqm/qmgrs'
```

then:

- 1) Take option 5 for **&SYSTEM** and check that the following directories are empty: isem, esem, msem, ssem, and shmem.
- 2) Take option 5 for **QMGRNAME** and check that the following directories are empty:- isem, esem, msem, ssem, and shmem.
- 3) Take option 5 for **&ipcc** in the QMGRNAME directory and check that the following directories are empty:- isem, esem, msem, ssem, and shmem.
- 4) Take option 5 for **&qmpersist** in the QMGRNAME directory and check that the following directories are empty:- isem, esem, msem, ssem, and shmem.
- 5) Take option 5 for **&app** and check that the following directories are empty: isem, esem, msem, ssem, and shmem.



## Save IBM MQ data on IBM i



Save IBM MQ data after removing unwanted FDC, trace, and JOB files.

### Before you begin

You need to have completed the tasks to remove unwanted and indoubt messages and quiesced IBM MQ.

### About this task

#### Procedure

1. Create a save file for every queue manager library on your system. To do so, issue the command:  
`CRTSAVF FILE(QGPL/ queue-manager-library )`  
  
where the *queue-manager-library* name consists of the name of the queue manager preceded by QM.
2. Save your queue manager libraries into the save files. To do so, issue the commands:  
`SAVLIB LIB( queue-manager-library ) DEV(*SAVF)`  
`SAVF(QGPL/ queue-manager-library )`
3. Remove all unwanted FDC data from directory:  
`QIBM/UserData/mqm/errors`
4. Remove old FDC files with the command:  
`RMVLNK OBJLNK('/QIBM/UserData/mqm/errors/*.FDC')`  
This command cleans up all files with an extension of 'FDC' in the IFS.
5. Remove old JOB files with the command:  
`RMVLNK OBJLNK('/QIBM/UserData/mqm/errors/*.JOB')`  
  
This command cleans up all files with an extension of 'JOB' in the IFS.
6. Remove all unwanted trace data from directory, or remove the whole directory:  
`QIBM/UserData/mqm/trace`
7. Remove all trace files with the command:  
`RMVLNK OBJLNK('/qibm/userdata/mqm/trace/*')`
8. Create a save file for IBM MQ IFS data. To do so, issue the command:  
`CRTSAVF FILE(QGPL/QMUSERDATA)`
9. Save your IBM MQ IFS data, using the command:  
`SAV DEV('/QSYS.LIB/QGPL.LIB/QMUSERDATA.FILE') OBJ('/QIBM/UserData/mqm')`
10. If you are going to run IBM MQ on a new machine, transfer the save files to the new machine.

## Install IBM MQ server on IBM i



Install the IBM MQ server in its primary language.

### Before you begin

You have completed planning the installation, obtained the installation disks, and set the system values; see Setting system values.

### About this task

Install the IBM MQ server and force object conversion. Object conversion migrates objects from the older to the newer version. By performing it now, rather than when an object is first used, you avoid slowing down the first use of the upgraded product.

After following the optional step to pre-agree the license, the **RSTLICPGM** command runs without requiring any interactive input. Otherwise the license agreement is displayed for you to accept. See License requirements.

### Procedure

1. Sign on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.
2. Optionally pre-agree the license terms and conditions by running the command,

```
CALL PGM (QSYS/QLPACAGR) PARM ('5724H72' 'V8R0M0' '0000' 0)
```

Where the parameters of **PARM** are,

#### **5724H72**

The product identifier for IBM MQ for IBM i.

#### **V8R0M0**

The version, release, and modification level.

#### **0000**

The option number for the \*BASE IBM MQ product option.

**0** Unused error structure.

3. Install IBM MQ for IBM i, base product, and primary language.

```
RSTLICPGM LICPGM (5724H72) DEV (install device) OPTION (*BASE) OUTPUT (*PRINT)
```

where the parameters of RSTLICPGM are,

#### **LICPGM ( 5724H72 )**

The product identifier for IBM MQ for IBM i.

#### **DEV ( *install device* )**

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

#### **OPTION ( \*BASE )**

Install the base IBM MQ for IBM i product.

#### **Unspecified parameters**

Unspecified parameters such as **RSTOBJ ( \*ALL )**, revert to defaults. The command installs both IBM MQ and the language files for the primary language of your system. For installing additional languages see Installing translated versions.

## What to do next

Install any Program Temporary Fixes (PTF) that have been issued.

## Install samples on IBM i



Install the IBM MQ samples

### Before you begin

If you have not already done so, sign on to the system with a user profile that has \*ALLOBJ special authority, for example QSEC0FR.

### About this task

Install the samples.

After following the optional step to pre-agree the license, the **RSTLICPGM** command runs without requiring any interactive input. Otherwise the license agreement is displayed for you to accept. See License requirements.

### Procedure

1. Optionally pre-agree the license terms and conditions by running the command,

```
CALL PGM (QSYS/QLPACAGR) PARM ('5724H72' 'V8R0M0' '0001' 0)
```

Where the parameters of **PARM** are,

**5724H72**

The product identifier for IBM MQ for IBM i.

**V8R0M0**

The version, release, and modification level.

**0001**

The option number for the samples.

**0** Unused error structure.

2. Install the samples using the command:

```
RSTLICPGM LICPGM (5724H72) DEV (install device) OPTION (1) OUTPUT (*PRINT)
```

Where the parameters of **RSTLICPGM** are,

**LICPGM ( 5724H72 )**

The product identifier for IBM MQ for IBM i.

**DEV ( install device )**

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

**OPTION ( 1 )**

Install the samples for IBM MQ for IBM i.

**OUTPUT ( \*PRINT )**

The output is printed with the spooled output of the job.

## Install translated versions on IBM i



Install translated versions of IBM MQ from a choice of national-languages.

### About this task

The following language versions are available for IBM MQ for IBM i:

Table 74. National-language versions of IBM MQ for IBM i

| Language ID | Language                                           |
|-------------|----------------------------------------------------|
| 2909        | Belgian English                                    |
| 2966        | Belgian French MNCS (Multi-National Character Set) |
| 2981        | Canadian French MNCS                               |
| 2975        | Czech                                              |
| 2950        | English uppercase                                  |
| 2924        | English uppercase and lowercase                    |
| 2984        | English US DBCS                                    |
| 2938        | English US uppercase DBCS                          |
| 2928        | French                                             |
| 2940        | French MNCS                                        |
| 2929        | German                                             |
| 2939        | German MNCS                                        |
| 2976        | Hungarian                                          |
| 2932        | Italian                                            |
| 2942        | Italian MNCS                                       |
| 2962        | Japanese                                           |
| 2986        | Korean                                             |
| 2978        | Polish                                             |
| 2979        | Russian                                            |
| 2989        | Simplified Chinese                                 |
| 2931        | Spanish                                            |

IBM MQ for IBM i is installed in the language that is the primary language on your system.

You can install additional versions of the product in any of the languages shown in Table 74. To do so:

### Procedure

1. Sign on to the system with a user profile that has \*ALLOBJ special authority
2. Issue the following command specifying the appropriate language ID:  

```
RSTLICPGM LICPGM(5724H72) DEV(install device) RSTOBJ(*LNG) LNG(language ID)
```

This installs the commands, message file, and panel groups into the relevant QSYS library for the language. For example, library QSYS2928 is used for French. If this QSYS29nn library does not exist, it is created by the RSTLICPGM command.

## Results

### Note:

1. To run the Japanese language version of IBM MQ for IBM i, the CCSID of the job must be 939 (5035) rather than 930 (5026) because IBM MQ uses lowercase English characters.
2. If you are installing IBM MQ for IBM i onto a machine for which the primary language is not on the CD, the install program prompts you to load a CD containing the product in that language. If, however, you have only one product CD, this means that the IBM MQ product has not been translated into your language. To get around this issue, proceed as follows:
  - Install the product in one of the supplied languages, and then add the corresponding QSYS29nn library into the *system library* list (for example using command CHGSYSLIBL). At the same time, check that there are no IBM MQ \*CMD, \*MENU, or \*MSGF objects in libraries higher up the library list. If some exist, then either delete these objects (because they refer to an earlier version of IBM MQ) or reorder the System Library list (because the product has been installed in more than one of the supplied languages).

## Verify the installation on IBM i



How to check that your installation has been successful.

### Procedure

1. To ensure that the product has loaded correctly, issue the Display Software Resources (DSPSFWRSC) command and check that the licensed program 5724H72 is listed. If you have installed the base and the optional samples, you see:

```
Resource
ID Option Feature Description
5724H72 *BASE 5050 IBM MQ for IBM i
5724H72 *BASE 2924 IBM MQ for IBM i
5724H72 1 5050 IBM MQ for IBM i - Samples
```

2. Press F11, while viewing the Display Software Resources screen, and you see the library and version number of the products installed:

```
Resource Feature
ID Option Feature Type Library Release
5724H72 *BASE 5050 *CODE QMQM V8R0M0
5724H72 *BASE 2924 *LNG QMQM V8R0M0
5724H72 1 5050 *CODE QMQMSAMP V8R0M0
```

3. If you have installed additional language versions, you also see entries for these versions. For example, if you have installed the French version, for which the language ID is 2928, you see:

a.

```
Resource
ID Option Feature Description
5724H72 *BASE 2928 IBM MQ for IBM i
```

b. and when you press F11:

```
Resource Feature
ID Option Feature Type Library Release
5724H72 *BASE 2928 *LNG QSYS2928 V8R0M0
```

4. Use the command DSPMQMVER to check exactly what version you have installed. For V8R0M0, it reports:

## Verify the upgrade on IBM i



After you have verified the installation, start the IBM MQ subsystem, check the queue managers, and take a fresh media recovery checkpoint.

### About this task

To verify that you have migrated to the latest version IBM MQ for IBM i, successfully:

### Procedure

1. Make QMQMADM either the primary or a secondary group profile for your user profile. To do so, issue one of the following commands:

```
CHGUSRPRF USRPRF(YOUR PROFILE) GRPPRF(QMQMADM)
CHGUSRPRF USRPRF(YOUR PROFILE) SUPGRPPRF(QMQMADM)
```

2. Start the IBM MQ subsystem with the command:

```
STRSBS SBS(QMQM/QMQM)
```

(If it is already running, you get error message CPF1010 which you can safely ignore).

3. Check that your queue managers are accessible by issuing the command:

```
WRKMQM
```

Use option 14 against each queue manager to start it.

Use option 5 against each queue manager to check its attributes.

4. You can use the other options to check your queue manager objects. For example, check your queues using option 18, check your channels using option 20, and so on.

5. Take a fresh media recovery checkpoint, using the following command:

```
RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(QMGRNAME) DSPJRNDTA(*YES)
```

Where *QMGRNAME* is the name of the queue manager.

## Restore queue managers after upgrading IBM MQ on IBM i



Complete the side-by-side upgrade by restored the saved queue managers onto the server that you have upgraded.

### Before you begin

**Note:** Carry out this task only if you are performing a side-by-side upgrade.

Ensure that you have saved your queue manager data, see “End IBM MQ activity on IBM i” on page 518, and installed and verified the upgrade.

### About this task

Transfer the queue manager data, and journal receivers, onto the server that has been upgraded.

### Procedure

1. Restore the queue manager libraries for every queue manager, using the command:

```
RSTLIB SAVLIB(queue-manager-library) DEV(*SAVF) (*PRINT)
SAVF(QGPL/ queue-manager-library)
```

where the *queue-manager-library* name consists of the name of the queue manager preceded by QM.

2. Restore the IBM MQ IFS data, using the command:  
RST DEV('/QSYS.LIB/QGPL.LIB/QMUSERDATA.FILE') OBJ('/QIBM/UserData/mqm') (\*PRINT)
3. To associate the journal receivers, issue the command WRKJRN on the journal AMQAJRN in each queue manager library, by pressing *PF4* and selecting option 9.
4. If you want to set up your work management environment, job descriptions, and pools, see the Administering IBMi for guidance. Otherwise, use the default setup.

## After Upgrading on IBM MQ for IBM i



Tasks to perform after you have upgraded IBM MQ for IBM i.

### About this task

Satisfy yourself the upgrade has completed successfully.

### Procedure

Delete the saved data in the save files in QGPL. This data was saved in “Save IBM MQ data on IBM i” on page 521.

## Post installation tasks for IBM MQ for IBM i



Tasks to perform after you have installed IBM MQ for IBM i, and before using it.

### About this task

When you have correctly installed IBM MQ for IBM i on your system:

### Procedure

1. For the latest product information for IBM i, see System requirements for IBM MQ .
2. To install and apply all fix packs, see “IBM i: Applying maintenance level updates on the latest release” on page 632.
3. Where you have more than one system and a mixture of releases of OS/400 or IBM i, and IBM MQ, you must take care when compiling CL programs. You must compile CL programs either on the system they are to run on, or on one with an identical combination of releases of OS/400 or IBM i, and IBM MQ. When you install later versions of IBM MQ, delete all IBM MQ commands from previous releases in any QSYSVvRrMm libraries using the QSYS/DLTCMD command.
4. If you have not installed IBM MQ on your system before, you must add user profiles to the QMQMADM group profile. Make all user profiles that are to be used for creating and administering queue managers members of the QMQMADM group profile, using the command CHGUSRPRF.
  - a. Start the IBM MQ subsystem, by issuing the command:  
STRSBS SBS(DQM/QMQM)

**Note:** The subsystem must be started after each IPL of the system, so you might choose to start it as part of your system startup process.

5. Create the system-default objects. The system-default objects are created automatically when you issue the CRTMQM command to create a queue manager. For example: CRTMQM MQMNAME(QMGRNAME)

ASP(\*SYSTEM). You can refresh them using the STRMQM command (Warning: this command will replace any existing default objects). For example: STRMQM MQMNAME(QMGRNAME) RDEFSYS(\*YES). Refer to the onscreen help for information about using this command.

**Note:** on the command STRMQM MQMNAME(QMGRNAME) RDEFSYS(\*YES):

- The command does not recreate the objects, it performs a CRTxxxx REPLACE(\*YES) for all of the SYSTEM.\* objects.
- This means that it refreshes the parameters on the objects back to their defaults. So if, for example, on the SYSTEM.DEFAULT.LOCAL.QUEUE object, TRGENBL had previously been changed to \*YES, then, when the command is run, it is changed back to TRGENBL(\*NO).
- If any messages exist on a queue, they are left intact, because the queues are not physically deleted.
- The contents of the SYSTEM.AUTH.DATA.QUEUE are untouched when this command is run.
- So, if the contents of this (or any other significant queue) become corrupt, it must be physically deleted and recreated either from scratch, or from a backup.

## Results

You are now ready to start using IBM MQ for IBM i.

**Note:** When you install IBM MQ for IBM i, two user profiles are created:

- QMQM
- QMQMADM

These two objects are central to the correct running of IBM MQ for IBM i. Do not alter or delete them. If you do, IBM cannot guarantee correct behavior of your product.

If you uninstall IBM MQ and data, these profiles are deleted. If you uninstall IBM MQ only, these profiles are retained.

## IBM i: migrating a queue manager from a previous release - alternative method



An alternative method of migrating a queue manager from a previous release to the latest release.

### Before you begin

1. Review the IBM MQ system requirements for the latest release of the product; see IBM MQ System Requirements
2. Review any other installed SupportPacs for their applicability to the latest release of IBM MQ.

### About this task

There are various parts to this form of migration:

1. As part of upgrading the IBM MQ product, carry out the following tasks:
  - a. "Preparing to install IBM MQ on IBM i" on page 529
  - b. "Install IBM MQ server on IBM i" on page 529
2. Following the IBM MQ product upgrade, carry out the following task:
  - a. "Post installation tasks" on page 531



## Preparing to install IBM MQ on IBM i



Carry out the following tasks to prepare your system for an upgrade.

### Procedure

1. Stop the IBM MQ queue managers by issuing the following command:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED) ENDCCTJOB(*YES) RCDMQMIMG(*YES)
TIMEOUT(30)
```

Ensure that the user profile issuing this command has \*ALLOBJ authority.

2. Create a save file for every queue manager library on your system. To do so, issue the command:

```
CRTSAVF FILE(QGPL/ queue-manager-library)
```

where the *queue-manager-library* name consists of the name of the queue manager preceded by QM.

3. Save your queue manager libraries into the save files. To do so, issue the commands:

```
SAVLIB LIB(queue-manager-library) DEV(*SAVF)
SAVF(QGPL/ queue-manager-library)
```

4. Create a save file for IBM MQ IFS data. To do so, issue the command:

```
CRTSAVF FILE(QGPL/QMUSERDATA)
```

5. Save your IBM MQ IFS data, using the command:

```
SAV DEV('/QSYS.LIB/QGPL.LIB/QMUSERDATA.FILE') OBJ('/QIBM/UserData/mqm')
```

6. If you are going to run IBM MQ on a new machine, transfer the save files to the new machine.
7. Issue the following command before you upgrade your IBM MQ product, only if the upgrade is required on the same machine.
  - a. `DLTMQM <QMGRNAME>`
  - b. `ENDSBS SBS(QMQM) OPTION(*IMMED)`
  - c. `WRKOBJLCK OBJ(QMQM) OBJTYPE(*LIB)`

Relinquish any locks on the system.

## Install IBM MQ server on IBM i



Install the IBM MQ server in its primary language and force object conversion.

### Before you begin

In either of the following cases, ensure that you have completed the planning and set the system values; see Setting system values

- If you have obtained the product through IBM passport advantage, follow the instructions in the EGA.README.txt file.
- If you have obtained the product on disk, follow the instructions within this topic.

### About this task

Install the IBM MQ server and force object conversion. Object conversion migrates objects from the older to the newer version. By performing it now, rather than when an object is first used, you avoid slowing down the first use of the upgraded product.

After following the optional step to pre-agree the license, the **RSTLICPGM** command runs without requiring any interactive input. Otherwise the license agreement is displayed for you to accept. See License requirements.

## Procedure

1. Sign on to the system with a user profile that has \*ALLOBJ special authority, for example QSECOFR.
2. Optionally pre-agree the license terms and conditions by running the command,

```
CALL PGM (QSYS/QLPACAGR) PARM ('5724H72' 'V8R0M0' '0000' 0)
```

Where the parameters of **PARM** are,

### **5724H72**

The product identifier for IBM MQ for IBM i.

### **V8R0M0**

The version, release, and modification level.

### **0000**

The option number for the \*BASE IBM MQ product option.

**0** Unused error structure.

3. Install IBM MQ for IBM i, base product, and primary language.

```
RSTLICPGM LICPGM (5724H72) DEV (install device) OPTION (*BASE) OUTPUT (*PRINT)
```

where the parameters of **RSTLICPGM** are,

### **LICPGM ( 5724H72 )**

The product identifier for IBM MQ for IBM i.

### **DEV ( install device )**

The device from which the product is to be loaded, typically an optical drive, for example, OPT01.

### **OPTION ( \*BASE )**

Install the base IBM MQ for IBM i product.

### **Unspecified parameters**

Unspecified parameters such as **RSTOBJ ( \*ALL )**, revert to defaults. The command installs both IBM MQ and the language files for the primary language of your system. For installing additional languages see Installing translated versions.

## What to do next

Install any Program Temporary Fixes (PTF) that have been issued.

To install the IBM MQ samples, see: "Install samples on IBM i" on page 523.

## Post installation tasks

Actions required after upgrading IBM MQ.

### About this task

Install the samples.

Carry out these steps after installing the product.

### Procedure

1. Issue the following commands:
  - a. STRSBS SBSD(QMQM/QMQM)
  - b. CRTMQM MQMNAME(<QMGRNAME>) DFTQMGR(\*YES) You receive the message " IBM MQ queue manager created."
  - c. STRMQM MQMNAME(<QMGRNAME>) You receive the message " IBM MQ queue manager '<QMGRNAME>' started."
2. Issue the following command:  
STRMQMMQSC SRCMBR(<QMGRNAME>) SRCFILE(\*CURLIB/QMQSC) OPTION(\*RUN)  
MQMNAME(<QMGRNAME>)
3. Reapply IBM MQ Authorities by issuing the command: CALL PGM(\*CURLIB/<QMGRNAME>)
  - a. You must compile the CLP as follows:  
CRTCLPGM PGM(\*CURLIB/<QMGRNAME>) SRCFILE(\*CURLIB/QMAUT) SRCMBR(\*PGM)

## Restoring a queue manager from the latest version to a previous version on UNIX systems and Windows

You can restore a queue manager to the previous version of the product from the latest version, if you have made a backup of the system or queue manager. If you have started the queue manager and processed any messages, or changed the configuration, the task cannot give you any guidance on restoring the current state of the queue manager.

### Before you begin

1. You must have made a backup of the system or queue manager before you upgraded to the latest version. For more information see Backing up and restoring IBM MQ queue manager data
2. If any messages were processed after starting the queue manager, you cannot easily undo the effects of processing the messages. You cannot restore the queue manager to the previous version of the product in its current state. The task cannot give you any guidance how to deal with subsequent changes that have occurred. For example, messages that were indoubt in a channel, or in a transmission queue on another queue manager, might have been processed. If the queue manager is part of a cluster, then configuration messages and application messages might have been exchanged.
3. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see setmqenv.

### About this task

When you restore a previous version of a queue manager, you restore the queue manager to its earlier code level. Queue manager data is restored to the state it was in when the queue manager was backed up.

### Procedure

1. Log in as a user in group mqm.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their SYSTEM.FTE.STATE queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.

- a. Run the **dspmq** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmq -o installation -o status
dspmq -a
```

**dspmq -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmq -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqlsr -m QMgrName
```

4. Restore the system, or IBM MQ and the queue manager.

If your backup procedure was to save the queue manager data, you must reinstall IBM MQ.

- a. Uninstall the previous installation.
- b. Reinstall IBM MQ from a manufacturing refresh.
- c. Apply the fix pack and interim fixes that restore IBM MQ to its previous level.
- d. Restore the queue manager data from the backup taken before installing the latest version.

5. Restart the previous version queue manager.

## What to do next

You might be restoring a previous version on a server with multiple IBM MQ installations. If one of the installations is primary, after restoring the previous version that installation, by default, becomes the primary installation.

You must review how applications connect to an installation. After restoring the previous version, some applications might connect to the wrong installation.

### Related information:

Backing up and restoring a queue manager

## Linux: Cleaning up after using the rpm freshen or upgrade options

The use of `rpm` upgrade or freshen options is not supported. If you use the options, follow this cleanup procedure, and then install following the correct steps.

### Before you begin

You have attempted to upgrade IBM MQ for Linux using `rpm -U` or `rpm -F`

### About this task

By using the freshen or upgrade options, you might have deleted your old IBM MQ package entries from the `rpm` database without removing the product from your system. You might also have partially installed IBM MQ

### Procedure

Follow these steps to clean up your system.

1. Find out which IBM MQ packages still have entries in your RPM database.

```
rpm -qa | grep MQSeries
```

2. Remove all remaining IBM MQ packages from your system.

```
rpm -e package-name
```

3. Remove the `/opt/mqm` directory.

```
rm -rf /opt/mqm
```

---

## Migrating IBM MQ library loading from an earlier version of the product to the latest version

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

### Before you begin

To migrate applications from an earlier version of the product to the latest version, you must know how the operating system loads a IBM MQ library for an application. Is the load path fixed by the application, and can you set the path in an environment variable? It is not essential to know the name of the IBM MQ library that the application loads. The library name does not change from an earlier version of the product to the latest version, although the contents of the library do.

## About this task

To migrate an application from an earlier version of the product to the latest version, you do not have to recompile or relink the application, because the IBM MQ libraries are compatible with later versions; see “Application compatibility and interoperability with later versions of IBM MQ” on page 467. You might have to configure the runtime environment differently, for the operating system to load the latest version of the IBM MQ library. If you replaced an earlier version of the product with the latest version, following the “Single-stage” approach; you do not need to do anything at all; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496<sup>11</sup>.

The latest version of the product provides two commands to assist you in configuring the runtime environment, **setmqinst** and **setmqenv**. **setmqinst** sets the primary installation; see **setmqinst**. **setmqenv** initializes the command environment by setting environment variables; see **setmqenv**.

## Procedure

How you configure the runtime environment depends on a number of factors, some of which apply to your configuration. Consider which of the following questions apply to your configuration.

1. Did you follow the build procedure documented in the product documentation for the earlier version of the product?

You might be following a different build procedure tailored to your development environment, or adapted from a development tool such as Microsoft Visual Studio.

2. How did you specify the load path for the earlier version?

3. Is the application is loaded by another environment, such as Eclipse, or an application server?

You must modify the parameters that govern how the parent environment loads applications, not the way the parent environment is loaded.

4. Is the configuration for Windows, or UNIX and Linux ?

On Windows, the functions performed by an application might require that the queue manager it connects to is associated with the primary installation.

5. What constraints and requirements do you have on how the load path is specified in the latest version?

Security rules might restrict the use of `LD_LIBRARY_PATH`.

6. Is the latest version of the product installed alongside the earlier version?

If Version 7.0.1 is installed:

- You cannot make a later installation primary.
- You cannot install the later version in the default installation path, that was referenced by applications in Version 7.0.1.

## What to do next

Environment configuration on the Windows, and UNIX platforms is a little different. If you have followed the documented build procedure in the earlier version, look at “Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 540 or “UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542. These topics show the effects of using **setmqinst** and **setmqenv** commands to configure the operating system environment for the three migration scenarios listed in Related information.

---

11. If you changed the location of the libraries from the earlier version, or created symbolic links to the libraries, this statement might not hold true.

**Related tasks:**

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** .

Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.


“UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542  
Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

“Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 540

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

**Related reference:**

“Coexistence” on page 581

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

**Related information:**

Changing the primary installation

Connecting applications in a multiple installation environment

setmqenv

setmqinst

setmqm

---

## Migrating an IBM MQ MQI client to the latest version of the product

Migrate an IBM MQ MQI client to the latest version by doing the tasks in the following topics. Check IBM MQ MQI client applications with the latest version before migration. You must stop all IBM MQ activity on the workstation, before upgrading the IBM MQ MQI client. After upgrading the IBM MQ MQI client, you must check the client channel configuration.

## Migrating an IBM MQ MQI client on UNIX systems, Windows, and IBM i to the latest version

Before migrating an IBM MQ MQI client, create a migration plan. Stop all IBM MQ activity on the client workstation. Upgrade the IBM MQ MQI client installation. Make any essential configuration and application changes.

### Before you begin

1. Create a migration plan. Use the planning task, “IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485, as a guide.

### Procedure

1. Review the IBM MQ system requirements for Version 8.0.  
See IBM MQ System Requirements.
2. Review all the changes in IBM MQ that affect you.  
See What's changed in IBM MQ Version 8.0.
3. End all IBM MQ activity on the workstation.
4. Upgrade the client. Select the appropriate platform that your enterprise uses.
  - To upgrade an IBM MQ MQI client for AIX installation on a workstation; see Client installation procedure on an AIX workstation.
  - To upgrade an IBM MQ MQI client for AIX installation on an AIX IBM MQ server; see Installing an IBM MQ MQI client on the same computer as the server.
  - To upgrade an IBM MQ MQI client for HP-UX installation on a workstation; see Client installation procedure on an HP-UX workstation.
  - To upgrade an IBM MQ MQI client for HP-UX installation on an HP-UX IBM MQ server; see Installing an IBM MQ MQI client on the same computer as the server.
  - To upgrade an IBM MQ MQI client for Linux installation on a workstation; see Client installation procedure on an Linux workstation.
  - To upgrade an IBM MQ MQI client for Linux installation on an Linux IBM MQ server; see Installing an IBM MQ MQI client on the same computer as the server.
  - To upgrade an IBM MQ MQI client for Solaris installation on a workstation; see Client installation procedure on an Solaris workstation.
  - To upgrade an IBM MQ MQI client for Solaris installation on an Solaris IBM MQ server; see Installing an IBM MQ MQI client on the same computer as the server.
  - To upgrade an IBM MQ MQI client for Windows installation on a workstation; see Client installation procedure on an Windows workstation.
  - To upgrade an IBM MQ MQI client for Windows installation on an Windows IBM MQ server; see Installing an IBM MQ MQI client on the same computer as the server.
  - To upgrade an IBM MQ MQI client for IBM i installation on a workstation; see Client installation procedure on IBM i .

### What to do next

Complete the tasks in your migration plan, such as verifying IBM MQ MQI client applications work correctly with the latest version.



**Related tasks:**

“IBM MQ migration planning to the latest version on UNIX platforms, Windows, and IBM i” on page 485  
Before migrating from one version to another, read this topic. Create your own migration plan based on the outline in the planning topic.

**Related information:**

Client installation procedure on an AIX workstation

Client installation procedure on an HP-UX workstation

Client installation procedure on a Linux workstation

Client installation procedure on a Solaris workstation

Client installation procedure on a Windows workstation

Client installation procedure on IBM i

Installing IBM MQ MQI clients on the same machine as the server

## Restoring an IBM MQ MQI client and client connection to the previous version

If you restore an IBM MQ MQI client from a later version of the product to an earlier version of the product, you must undo the configuration changes manually.

### About this task

It is unusual to restore earlier IBM MQ MQI client libraries to a workstation. The principal tasks are listed in the following steps.

### Procedure

1. End all IBM MQ activity on the workstation.
2. Uninstall the later version of the IBM MQ MQI client code.
3. Follow the client installation procedure for the platform to install the earlier version of the IBM MQ MQI client code.
4. If you configured a Client Connection Definition Table (CCDT) for a queue manager on a later version of the product, revert to using a table created by a queue manager on the earlier version.

The CCDT must always be created by a queue manager on the same, or earlier, release to the client.

---

## Migrating applications to the latest version of the product

IBM MQ applications might require migration between Version 7.1 and the latest version.

### About this task

For more information see [Connection authentication: Configuration](#).

For a new IBM MQ Version 8.0 installation, the **CONNAUTH CHCKLOCL** attribute will be set to **OPTIONAL**. This means that user IDs and passwords are not required, but if they are provided they must be a valid pair, or they will be rejected.

When you are migrating between IBM WebSphere MQ Version 7.1 and latest version, the **CONNAUTH CHCKLOCL** attribute on each queue manager is set to **NONE**, ensuring version to version continuity, but switching connection authentication off.

## Migrating IBM MQ library loading from an earlier version of the product to the latest version

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

### Before you begin

To migrate applications from an earlier version of the product to the latest version, you must know how the operating system loads a IBM MQ library for an application. Is the load path fixed by the application, and can you set the path in an environment variable? It is not essential to know the name of the IBM MQ library that the application loads. The library name does not change from an earlier version of the product to the latest version, although the contents of the library do.

### About this task

To migrate an application from an earlier version of the product to the latest version, you do not have to recompile or relink the application, because the IBM MQ libraries are compatible with later versions; see “Application compatibility and interoperability with later versions of IBM MQ” on page 467. You might have to configure the runtime environment differently, for the operating system to load the latest version of the IBM MQ library. If you replaced an earlier version of the product with the latest version, following the “Single-stage” approach; you do not need to do anything at all; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496<sup>12</sup>.

The latest version of the product provides two commands to assist you in configuring the runtime environment, **setmqinst** and **setmqenv**. **setmqinst** sets the primary installation; see **setmqinst**. **setmqenv** initializes the command environment by setting environment variables; see **setmqenv**.

### Procedure

How you configure the runtime environment depends on a number of factors, some of which apply to your configuration. Consider which of the following questions apply to your configuration.

1. Did you follow the build procedure documented in the product documentation for the earlier version of the product?  
You might be following a different build procedure tailored to your development environment, or adapted from a development tool such as Microsoft Visual Studio.
2. How did you specify the load path for the earlier version?
3. Is the application is loaded by another environment, such as Eclipse, or an application server?  
You must modify the parameters that govern how the parent environment loads applications, not the way the parent environment is loaded.
4. Is the configuration for Windows, or UNIX and Linux ?  
On Windows, the functions performed by an application might require that the queue manager it connects to is associated with the primary installation.
5. What constraints and requirements do you have on how the load path is specified in the latest version?  
Security rules might restrict the use of LD\_LIBRARY\_PATH.

---

12. If you changed the location of the libraries from the earlier version, or created symbolic links to the libraries, this statement might not hold true.

6. Is the latest version of the product installed alongside the earlier version?

If Version 7.0.1 is installed:

- You cannot make a later installation primary.
- You cannot install the later version in the default installation path, that was referenced by applications in Version 7.0.1.

## What to do next

Environment configuration on the Windows, and UNIX platforms is a little different. If you have followed the documented build procedure in the earlier version, look at “Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 540 or “UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542. These topics show the effects of using `setmqinst` and `setmqenv` commands to configure the operating system environment for the three migration scenarios listed in Related information.

### Related tasks:

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.


“UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542  
Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

“Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 540

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

### Related reference:

“Coexistence” on page 581

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

### Related information:

Changing the primary installation  
 Connecting applications in a multiple installation environment  
 setmqenv  
 setmqinst  
 setmqm

## Windows: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

### Before you begin

Read “Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460 and “Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533 before starting this task.

Plan and install the latest version of IBM MQ for Windows, and remember the installation name and whether the installation was set to primary.

### About this task

Windows searches numerous directories for load libraries, called DLL s; see Dynamic-Link Library Search Order.

The build procedure documented for IBM WebSphere MQ Version 7.0.1 applications is to place the IBM MQ libraries to load before any other product libraries in the `cl` command. The IBM MQ `.lib` libraries must be in the `PATH` environment variable you have specified at build time, and the `DLL` libraries at run time. The `PATH` variable is used by the application process to find the libraries it must load. If you have followed this build procedure, then the effect of installing the latest version of the product on the libraries that are loaded depends on the migration scenario; see Table 75.

Table 75. Windows configurations

|                                | Scenario | Latest version replaces earlier version in the same location                                                                                                                                                                                                                                                             | Latest version replaces earlier version in a different location                                                                                                                                                                                                                                                                       | Latest version alongside earlier version<br>“Multi-stage”                                                        |
|--------------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Action                         |          | “Single-stage”                                                                                                                                                                                                                                                                                                           | “Side-by-side”                                                                                                                                                                                                                                                                                                                        |                                                                                                                  |
| <b>setmqinst</b>               |          | <b>setmqinst</b> makes the latest version installation primary. The global <code>PATH</code> is changed to point to the latest version library and all Windows features work with the latest version <sup>See note</sup> .                                                                                               |                                                                                                                                                                                                                                                                                                                                       | No. The latest version installation can be primary, because an earlier version is installed.                     |
| No other configuration actions |          | Library loading works correctly.<br><br>The global <code>PATH</code> contains the location of the latest version libraries.<br><br>Even if the latest version installation is not primary, library loading works correctly. The latest version libraries are in the same location as the earlier version libraries were. | Library loading probably works correctly.<br><br>The library loading might not work, if the application process locally modified the <code>PATH</code> to reference the location of the earlier version libraries. A local setting of <code>PATH</code> might override the global <code>PATH</code> that is set by <b>setmqinst</b> . | The library loading continues to work with the earlier version correctly, nothing works with the latest version. |

Table 75. Windows configurations (continued)

|                 | Scenario | Latest version replaces earlier version in the same location<br>“Single-stage”         | Latest version replaces earlier version in a different location<br>“Side-by-side” | Latest version alongside earlier version<br>“Multi-stage”                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------|----------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action          |          |                                                                                        |                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>setmqenv</b> |          | Library loading works correctly.<br><br><b>setmqenv</b> sets the local PATH correctly. |                                                                                   | Library loading works correctly, both for the earlier version and the latest version.<br><br><b>setmqenv</b> sets the local PATH correctly for the latest version. But the Windows features that depend on the global path do not work correctly with the latest version <sup>See note</sup> .<br><br>The correct earlier version is loaded, because the latest version library loads the earlier version library for queue managers that have not been migrated from the earlier version. |

## Procedure

Identify the installation of the latest version of the product, from which the operating system is going to load IBM MQ libraries:

- If you have a multiple installations of the latest versions to load from a server, IBM MQ checks that the installation the library was loaded from is the installation that is associated with any queue manager the application calls. IBM MQ loads the correct library if the wrong library is loaded. It is necessary to configure only one runtime environment for all IBM MQ applications.
- A typical choice is set the primary installation. Setting an installation to be primary places its library path in the global PATH variable.
- If you upgraded an earlier version installation to the latest version, a link path to the earlier version installation now points to an installation containing the latest version. Applications that have a fixed linkage path to the earlier version installation now load the libraries for the latest installation. They are then switched to the installation that is associated with any queue manager they connect to.
- If you rebuild an application, it must link to an installation of the latest version.
- If an application uses COM or ActiveX it can connect to any queue manager as long as there is a primary installation and it is Version 7.1 or later.

**Note:** If an earlier version of the product is installed, COM or ActiveX server applications connect to queue managers associated only with the Version 7.0.1 installation. COM or ActiveX client applications are not affected by the limitation.

- If you are running the IBM MQ.NET monitor in transactional mode, the queue manager it connects to must be the primary installation.

## What to do next

If you add further installations of the latest version of the product, you must decide which installation to make primary, if you have chosen to make any primary. As long as applications load IBM MQ libraries from one of the latest version installations, such as the primary installation, they can connect to queue managers associated with any other latest version installation.

On Windows, you might build applications with different development tools. You must identify the property of the development tool that sets the PATH of the application that is being built, and not the properties of the tool itself. For example, if you are debugging with Microsoft Visual Studio, you can insert a call to `setmqenv` in the **Environment** property of the debugging section of the **Configuration** properties of a project.

A Windows application might call LoadLibrary and specify an explicit load path. You might build a side-by-side assembly and configure an explicit load path. If an application uses either of these mechanisms, and the latest version IBM MQ library is not on the same path as the earlier release, you must recompile, or configure and relink your application to load the latest version libraries.

#### **Related tasks:**

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500


Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

#### **Related reference:**

“Coexistence” on page 581

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

#### **Related information:**

Changing the primary installation

Connecting applications in a multiple installation environment

`setmqenv`

`setmqinst`

`setmqm`

Features that can be used only with the primary installation on Windows

## **UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version**

Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

## Before you begin

Read “Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460 and “Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533 before starting this task.

Plan and install the latest version of IBM MQ for Windows, and remember the installation name and whether the installation was set to primary.

## About this task

In Version 7.0.1, the documented build procedure for IBM MQ applications is to include an explicit library path to the location of the IBM MQ libraries, and to `/usr/lib`, in the link step of the compiler; see Figure 66. The same build procedure is documented for the latest version of the product.

---

```
gcc -m32 -o amqsput_32_r amqsput0.c -I/opt/mqm/inc -L/opt/mqm/lib
-Wl,-rpath=/opt/mqm/lib -Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

---

Figure 66. Linux C server application, 32 bit, threaded compile and link Version 7.0.1

The documented build step for other UNIX platforms is similar. The examples in Table 76 are all based on Linux.

If you have followed this build procedure, then the effect of installing the latest version on library loading depends on the migration scenario; see Table 76:

Table 76. UNIX and Linux configurations

|                                | Scenario | Latest version replaces earlier version in the same location<br><br>“Single-stage”                                                                                                                                                                                                    | Latest version replaces earlier version in a different location<br><br>“Side-by-side”                                                                                                  | Latest version alongside earlier version<br><br>“Multi-stage”                                                    |
|--------------------------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Action                         |          |                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                        |                                                                                                                  |
| <b>setmqinst</b>               |          | <b>setmqinst</b> makes the latest version installation primary. Symbolic links to the IBM MQ link libraries are inserted into <code>/usr/lib</code> .                                                                                                                                 |                                                                                                                                                                                        | No. The latest version installation can be primary, because an earlier version is installed.                     |
| No other configuration actions |          | Library loading works correctly.<br><br>Library loading works, even without the latest version installation being made primary, because the libraries are installed in <code>/opt/mqm/lib</code> and the application was built with the link option, <code>-rpath=/opt/mqm/lib</code> | Library loading works correctly.<br><br>Library loading works, because the installation is primary, and the application was built with the link option, <code>-rpath=/usr/lib</code> . | The library loading continues to work with the earlier version correctly, nothing works with the latest version. |

Table 76. UNIX and Linux configurations (continued)

|                                                         | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Latest version replaces earlier version in the same location<br><br>"Single-stage"                                                                                                                       | Latest version replaces earlier version in a different location<br><br>"Side-by-side"                                                                                                                                                                                             | Latest version alongside earlier version<br><br>"Multi-stage" |
|---------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|
| Action                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                   |                                                               |
| <b>setmqenv</b> , without setting the -k or -l options. | Library loading works correctly.<br><br><b>setmqenv</b> is unnecessary. Library loading works, because the libraries are installed in /opt/mqm/lib and the application was built with the link option, -rpath=/opt/mqm/lib.                                                                                                                                                                                                                                                 | Library loading works correctly.<br><br><b>setmqenv</b> is unnecessary. Library loading works, because the installation is primary, and the application was built with the link option, -rpath=/usr/lib. | The library loading continues to work with the earlier version correctly, nothing works with the latest version.                                                                                                                                                                  |                                                               |
| <b>setmqenv</b> , with the -k or -l options set.        | Library loading works correctly.                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                          | Library loading works correctly, both for the earlier version and the latest version.<br><br>The correct earlier version is loaded, because the latest version library loads the earlier version library for queue managers that have not been migrated from the earlier version. |                                                               |
|                                                         | The operating system finds the IBM MQ library location set by <b>setmqenv</b> . <b>setmqenv</b> adds the location to LD_LIBRARY_PATH ( LIBPATH on AIX. On HP-UX LD_LIBRARY_PATH is set, not SHLIB_PATH.) . LD_LIBRARY_PATH is searched before paths set in the application or paths in the default search path. Not all applications can load a library using LD_LIBRARY_PATH. In which case the application works only if the library location is /opt/mqm/lib or /usr/lib |                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                   |                                                               |

## Procedure

Identify the installation of the latest version of the product, from which the operating system is going to load IBM MQ libraries:

- If you have a multiple installations of the latest versions to load from a server, IBM MQ checks that the installation the library was loaded from is the installation that is associated with any queue manager the application calls. IBM MQ loads the correct library if the wrong library is loaded. It is necessary to configure only one runtime environment for all IBM MQ applications.
- A typical choice is to set the primary installation. Setting an installation to be primary places symbolic links to the IBM MQ libraries in /usr/lib. Applications built following the Version 7.0 instructions have an explicit link to /usr/lib. /usr/lib is also normally in the default library search path.
- If you upgraded an earlier version installation to the latest version, a link path to the earlier version installation now points to an installation containing the latest version. Applications that have a fixed linkage path to the earlier version installation now load the libraries for the latest installation. They are then switched to the installation that is associated with any queue manager they connect to.
- If you rebuild an application, it must link to an installation of the latest version.
- If you set LD\_LIBRARY\_PATH, or LIBPATH on AIX, you must check that the application is able to use LD\_LIBRARY\_PATH. setuid or setgid, applications, or applications built in other ways, might ignore LD\_LIBRARY\_PATH for security reasons.



## What to do next

If you add further installations of the latest version of the product, you must decide which installation to make primary, if you have chosen to make any primary. As long as applications load IBM MQ libraries from one of the latest version installations, such as the primary installation, they can connect to queue managers associated with any other latest version installation.

### Related tasks:

“UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496

Single-stage migration, is the term used to describe replacing the only installation of IBM MQ on a server, with a later release. Single stage migration is also known as **upgrading in place** or **in place upgrade** . Until Version 7.0.1.6, single-stage was the only migration scenario. Single-stage migration preserves existing scripts and procedures for running IBM MQ the most. With other migration scenarios you might change some scripts and procedures, but you can reduce the effect queue manager migration has on users.

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500


Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

### Related reference:

“Coexistence” on page 581

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

### Related information:

External library and control command links to primary installation on UNIX and Linux

Connecting applications in a multiple installation environment

Changing the primary installation

setmqenv

setmqinst

setmqm

Loading IBM WebSphere MQ Version 7.1 libraries

## Linux: Rebuilding a C++ application

C++ IBM MQ MQI client and server applications on Linux must be recompiled using GNU Compiler Collection (GCC) 4.1.2, or later. Compilers older than GCC 4.1.2 are no longer supported. The C++ GCC 4.1.2 run time libraries, or later, must be installed in `/usr/lib` or `/usr/lib64`

If you are using one of the supported Linux distributions, the libraries are correctly installed; see IBM MQ System Requirements.

The GCC 4.1.2 libraries support SSL and TLS connections from an IBM MQ MQI client. SSL and TLS use GSKit version 8, which depends on `libstdc++.so.6`. `libstdc++.so.6` is included in GCC 4.1.2.

## Before you begin

1. Check the required level of GCC for your distribution of Linux ; see IBM MQ System Requirements.
2. If you are using SSL or TLS, also check the required level of `libstdc++.so`.
3. Check whether the application requires rebuilding. Run the following command to display what version of `libstdc++.so` the application depends upon. If the result is less than `libstdc++.so.6`, you must rebuild your application.

```
ldd ApplicationPath
```

## About this task

The task describes the steps required to rebuild a Linux C++ IBM MQ application. For more detailed instructions about building Linux applications for IBM MQ ; see Building your procedural application on Linux

## Procedure

1. Check that the required GCC library is installed correctly.

Run one of the following commands:

- Check the 32 bit library on an x86 Linux system:

```
ls -l /usr/lib/libstdc++.so.6
```

- Check the 64 bit library on any other Linux system.

```
ls -l /usr/lib64/libstdc++.so.6
```

2. Check that the GCC compiler is at least at version 4.1.2

Run the following command to display the version of GCC.

```
gcc -v
```

3. Rebuild the application

The commands to compile and link Linux C++ applications are described in Building 32-bit applications and Building 64-bit applications

## What to do next

When you deploy your Linux C++ application, ensure that the same GCC runtime library is correctly installed on the run time system.

---

## Migrating IBM MQ for z/OS - order of tasks

Perform these instructions, in the order shown, to migrate a single IBM MQ for z/OS queue manager.

### About this task

The tables within this topic show the tasks required in each part of the process to migrate IBM MQ for z/OS, and the order in which these tasks must be done.

#### Notes:

- You must perform the tasks in the following order:

1. Before migration
2. Migrating to the next release
3. Post-migration tasks

and the order of tasks within each table.

*Table 77. Before migration*

| Task                                              | For your own use |
|---------------------------------------------------|------------------|
| 1. Make your configuration ready for migration    |                  |
| 2. Install the new version code                   |                  |
| 3. Perform backup                                 |                  |
| 4. Restart IBM MQ systems                         |                  |
| 5. Review pageset 0 usage                         |                  |
| 6. Migrate Db2 tables - for each QSG              |                  |
| 7. Add new CF structure definition - for each QSG |                  |
| 8. Server application migration                   |                  |
| 9. Preparing to migrate Advanced Message Security |                  |

*Table 78. Migrating to the next release*

| Task                                                           | For your own use |
|----------------------------------------------------------------|------------------|
| 10. Stop or disconnect all your applications                   |                  |
| 11. Stop the queue manager                                     |                  |
| 12. Update STEPLIB for MSTR and channel initiator              |                  |
| 13. Update the initialization input data sets                  |                  |
| 14. Update the target version system parameter module          |                  |
| 15. Migrating Advanced Message Security                        |                  |
| 16. Review the security control of your system                 |                  |
| 17. Start the queue manager                                    |                  |
| 18. Optionally, revert the queue manager to a previous version |                  |

Table 79. Post migration tasks

| Task                                                   | For your own use |
|--------------------------------------------------------|------------------|
| 19. Check the changes in behavior                      |                  |
| 20. Modify the backup jobs                             |                  |
| 21. Post migration tasks for Advanced Message Security |                  |
| 22. Perform a full regression test                     |                  |
| 23. Update the ZPARM module, if not already done       |                  |
| 24. Set OPMODE to NEWFUNC                              |                  |
| 25. Exploit the new function                           |                  |
| 26. Consider client application migration              |                  |

## z/OS: new messages in IBM MQ for z/OS Version 8

Read this information to see the list of new messages in IBM MQ for z/OS Version 8

### About this task

#### Recovery log manager messages

CSQJ164I *csect-name* Log archiving delayed, all available offload tasks in use

CSQJ168I *csect-name* Log archiving is no longer delayed

#### Message manager messages

CSQM079I *csect-name* Policy access attempt rejected due to incompatible AMS version, jobname *jobname*

CSQM523I *csect-name* CLUSTER OR CLROUTE CANNOT CURRENTLY BE ALTERED

CSQM526I *csect-name* CLUSTER OR CLROUTE CANNOT CURRENTLY BE ALTERED *csect-name*  
CERTIFICATE LABEL NOT ALLOWED FOR SSL 3.0 CHANNEL

#### Buffer manager messages

CSQP054I Buffer pool *n* is now located above the bar

CSQP055I Buffer pool *n* is now located below the bar

CSQP056E The ALTER BUFFPOOL command for buffer pool *n* has failed

#### Topic manager messages

CSQT824I *csect-name* Topic *topic-1* is dependent on PROXYSUB(FORCE) of topic *topic-2* from a different Pub/Sub hierarchy stream

CSQT967E *csect-name* Unable to deliver proxy subscription to queue manager *queue-manager*, reason=*mqrc (mqrc-text)*

CSQT968I *csect-name* Topic *topic-1* in cluster *cluster-name* is dependent on PROXYSUB(FORCE) of topic *topic-2*

CSQT971E *csect-name task* failed to quiesce

CSQT983E *csect-name task* failed, reason *mqrc (mqrc-text)*, retry in *n* minutes

CSQT984E *csect-name task* has encountered *n* occurrences of reason *mqrc (mqrc-text)* while attempting to process a message.

CSQT987E *csect-name task* failed due to reason *mqrc (mqrc-text)* Retry in *n* minutes

CSQT988E *csect-name task* failed due to reason *mqrc (mqrc-text)* Retry in *n* minutes

CSQT989E *csect-name task* has encountered *n* occurrences of reason *mqrc (mqrc-text)* while attempting to process a message.

CSQT990E *csect-name task* has encountered *n* occurrences of reason *mqrc (mqrc-text)* while attempting to process a message.

CSQT991I *csect-name task* has recovered from previous error condition

CSQT996E *csect-name* Creation of proxy subscription failed on queue manager *qmgr-name*, cluster *cluster-name*, topic string *topic-string*, reason=*mqrc (mqrc-text)*

CSQT997E *csect-name* Cancellation of proxy subscription failed on queue manager *qmgr-name*, cluster *cluster-name*, topic string *topic-string*, reason=*mqrc (mqrc-text)*

CSQT998E *csect-name* Proxy subscription re-synchronization failed on queue manager *qmgr-name*, cluster *cluster-name*, reason=*mqrc (mqrc-text)*

CSQT999E *csect-name* Proxy subscription re-synchronization failed on queue manager *qmgr-name*, cluster *cluster-name*, reason=*mqrc (mqrc-text)**csect-name task* has encountered a message that is not valid on queue *queue*

#### Utilities messages

CSQU179E The transmission queue cannot be switched because the channel initiator is not active

#### Distributed queuing messages

CSQX469E *csect-name* Update not received for CLUSRCVR channel *channel-name* hosted on queue manager *qmgr* in cluster *cluster-name*, expected *n* days ago, *m* days remaining

CSQX878I *csect-name* Repository command error, command *command*, cluster object *object-name*, sender *sender-id*, reason *reason*

CSQX879E *csect-name* Conflicting clustered topic *topic-name* from queue manager *qmgr-name*

#### Initialization procedure and general services messages

CSQY024I IBM MQ AMS for z/OS is not installed, but the system parameter SPLCAP is set to YES

CSQY025I IBM MQ AMS for z/OS is not installed, but the system parameter SPLCAP is set to YES  
IBM MQ AMS for z/OS is installed.

CSQY336E *csect-name keyword* not allowed - restricted functionality

CSQY337E *csect-name keyword* value length not allowed - restricted functionality

#### IBM MQ Advanced Message Security messages

CSQ0417I Quality of protection: *qop*

CSQ0418I Toleration: *toleration-flag*

CSQ0468I No policies found

#### Service facilities messages

CSQ1134E KEYWORD EXTRACT REQUIRES AT LEAST ONE OUTPUT DDNAME

CSQ1219I LOG RECORDS CONTAIN *n* BYTE RBA - QSG(*in-qsg*)

## z/OS: overall migration - order of tasks

Read this information that shows an overall migration plan for your system, together with the order in which you must undertake the tasks.

### Before you begin

Read the information in Clustering: Best practices and Clustering: Topology design considerations to understand repositories.

### About this task

The tables within this topic show the tasks required in each part of the process to migrate an overall system, and the order in which these tasks must be done.

#### Notes:

- You must perform the tasks in the following order:
  1. Overview of migration
  2. Migrating your system to the next release

Queue managers in a queue sharing group (QSG) and queue managers in a cluster can be migrated in parallel, but at any time there should be enough queue managers working in the QSG, and cluster, to ensure that your business can operate satisfactorily while a staged migration takes place.

If there are queue managers in clusters, ensure that you migrate the full repository queue managers before migrating any partial repository queue managers, so that the cluster always has a full picture of the current cluster architecture.

*Table 80. Overview of migration*

| Task                                           | For your own use |
|------------------------------------------------|------------------|
| 1. Make your configuration ready for migration |                  |
| 2. Install the new code                        |                  |
| 3. Server application migration                |                  |

*Table 81. Migrating your system to the next release*

| Task                                                                    | For your own use |
|-------------------------------------------------------------------------|------------------|
| Full repositories                                                       |                  |
| Overview - Migrate one full repository on IBM MQ for z/OS in a QSG      |                  |
| 4. Migrate DB/2 tables and New CF structure definitions                 |                  |
| 5. Migrate the queue manager                                            |                  |
| 6. Migrate other full repositories on IBM MQ for z/OS, if any           |                  |
| 7. Full repositories not in a QSG                                       |                  |
| 8. Full repositories on other platforms                                 |                  |
| Other queue managers                                                    |                  |
| Overview - Migrate partial repositories on IBM MQ for z/OS, in a QSG    |                  |
| 9. Migrate DB/2 tables and New CF structure definitions if not yet done |                  |
| 10. Migrate one queue manager                                           |                  |

Table 81. Migrating your system to the next release (continued)

| Task                                                                   | For your own use |
|------------------------------------------------------------------------|------------------|
| 11. Test the migration and upgrade the other queue managers            |                  |
| 12. Migrate the partial repositories not in a QSG                      |                  |
| Migrate queue managers not in a QSG or cluster on IBM MQ for z/OS      |                  |
| 13. Migrate partial repositories on other platforms                    |                  |
| 14. Migrate queue managers not in a cluster on other platforms         |                  |
| Testing and using new function; see "Post migration tasks" on page 564 |                  |
| 15. Regression tests                                                   |                  |
| 16. Use the new function available to you                              |                  |

You can upgrade the client libraries at any time during the process. As a final task, recompile the clients using new functions and deploy.

## Results

You have migrated your system to another release.

## Migrating from earlier unsupported releases of IBM MQ for z/OS

You must take into consideration whether you are upgrading a production system, or a test system, before you undertake the migration process.

### Production systems

For production systems you should, firstly, migrate the unsupported release of IBM MQ to IBM WebSphere MQ Version 7.1, following the instructions given in the documentation for that release. For more information about where to find the documentation for older versions of the product, see Documentation for older versions of IBM MQ

You can then migrate to IBM MQ Version 8.0 following the instructions in this section.

**Important:** Ensure that your system is stable at Version 7.1, before migrating to Version 8.0, so that you have a system to revert to, if necessary.

### Test systems

For a test system, it might be appropriate to migrate directly to IBM MQ Version 8.0 instead.

You should take complete backups of your system, to ensure that you can restart from backups if you need to use the old release again.

IBM MQ Version 8.0 will migrate IBM MQ objects and messages during the first startup at Version 8.0.

New attributes added to objects in IBM WebSphere MQ Version 7.0 and IBM MQ Version 8.0 releases will be set to their default values.

After you have migrated to IBM MQ Version 8.0 using this method, you cannot revert to the original version.

You can restart a queue manager, at the original version, using the full set of backups taken before migration.

Note, that any changes you make to the system after the backups were taken, or while running at Version 8.0, will be lost.

## Backward migration to earlier supported releases of IBM MQ for z/OS

After installation of a new release of IBM MQ for z/OS, you carry out queue manager migration by stopping the queue manager, which is running with the prior release of code, and restarting the queue manager using the new release of code.

### Maintenance in a queue sharing group

In a queue sharing group, individual queue managers can be migrated forwards to IBM MQ Version 8.0.0, while those that remain at either IBM WebSphere MQ Version 7.0.1 or Version 7.1.0 can continue to function. This allows you to upgrade queue-sharing group queue managers to Version 8.0.0 at different times, maintaining the high availability of the queue sharing group.

The function required to enable lower level queue managers to tolerate Version 8.0.0 additions to QSGDISP(GROUP) and QSGDISP(SHARED) objects is incorporated in the same authorized program analysis reports (APARs) which provide backward migration capability.

### Code levels supported

Migration support is provided from both IBM WebSphere MQ Version 7.0.1 and IBM WebSphere MQ Version 7.1.0 to IBM MQ for z/OS Version 8.0.0.

The backward migration APAR for both IBM WebSphere MQ Version 7.0.1 and IBM WebSphere MQ Version 7.1.0 is PI19721.

**Important:** PTFs for this APAR must be applied on Version 7.0.1 or IBM WebSphere MQ Version 7.1.0 prior to attempting to fall back from IBM MQ for z/OS Version 8.0.0.

PTFs for this APAR are the *Migration and Toleration PTFs for Version 8.0.0* described in Planning for migration to the latest release.

Service has been discontinued for versions of the product prior to IBM WebSphere MQ Version 7.0.1. No backward migration capability is available for these versions.

The IBM MQ for z/OS Version 8.0.0 early code installed in the link pack area (LPA) is downward compatible. The code supports queue managers running at Version 7.0.1 and Version 7.1.0.

Once updated to the Version 8.0.0 level, and the queue manager subsystem refreshed using the REFRESH QMGR TYPE(EARLY) command, the early code need not be changed for any subsequent forward or backward migration activity

Message

```
CSQ3111I <cpf> CSQYSCMD - EARLY PROCESSING PROGRAM IS V8.0.0 LEVEL 007-001
```

is displayed during startup in the queue manager joblog and indicates that the queue manager is using the correct level of early code.



## Limitations and restrictions

IBM MQ for z/OS Version 8.0.0. uses a migration switch to support backward migration by preventing use of certain new functions, that cannot be backward migrated, until the installation confirms that backward migration is no longer required.

The migration switch is configured through a change to ZPARM using the OPMODE parameter of CSQ6SYSP.

While OPMODE is set to COMPAT, it is possible to backward migrate, although certain new functions are not available. Once OPMODE is set to NEWFUNC, all new functions are available, but it is no longer possible to perform backward migration

The MQSC command DISPLAY SYSTEM shows the current value of OPMODE. The command displays two values, the operation mode, either COMPAT or NEWFUNC, and a version number.

When the operation mode is COMPAT, then the version number indicates to which version of IBM MQ for z/OS you can fall back.

*The value of OPMODE displayed during startup in message CSQY101I reflects the operation mode requested using ZPARM. Queue manager initialization evaluates the requested operation mode in combination with local state and other members of the queue sharing group, to determine the actual operation mode displayed with DISPLAY SYSTEM.*

You cannot backward migrate a queue manager, newly created at Version 8.0.0, to a prior release. A queue manager migrated forward to Version 8.0.0 *remembers* where it was migrated from, and it is only possible to fall back to that *remembered* prior version.

Certain connection types (IMS, BATCH and RRSBATCH used by WAS and Db2 stored procedures) allow an application to connect to multiple queue managers concurrently. If required, these queue managers can be running different levels of IBM MQ code. In such a scenario, the adapter code (usually referenced through a STEPLIB DD statement or environment variable) must be loaded from libraries corresponding with the highest level of the queue managers connected. This ability for the adapter code to support connections to older queue managers means that in a backward migration scenario it is possible to just restart the MSTR and CHIN procedures with the back level code, and not change connecting jobs.

The operations and controls ISPF panels, CSQOREXX, from IBM MQ for z/OS Version 8.0.0 , are able to connect to and administer queue managers from a prior release. However, the ISPF panels from lower releases are not able to connect to IBM MQ for z/OS Version 8.0.0. When migrating, or during fall back, either use the same version ISPF panels as the level of code the queue manager is running, or use CSQOREXX from the higher release of code. In a mixed level queue sharing group, the IBM MQ for z/OS Version 8.0.0 panels must be used to administer Version 7.1.0 or Version 7.0.1 queue managers, as ISPF panels from earlier releases do not tolerate responses from any Version 8.0.0 queue managers.

### Queue sharing groups

Prior to migrating any queue manager in a queue sharing group to IBM MQ for z/OS Version 8.0.0 you must ensure that all the queue managers in the queue sharing group are at the same version and release. That is, they must all be at IBM WebSphere MQ Version 7.0.1 or, all at IBM WebSphere MQ Version 7.1.0.

**Attention:** Once you have started the migration process to IBM MQ for z/OS Version 8.0.0 for queue managers in a queue sharing group, it is not possible to backward migrate any other queue manager to an earlier release. This would lead to three different versions in the queue sharing group, which is not permitted.

### Related reference:

“z/OS: Switching from OPMODE=(NEWFUNC,800) to OPMODE=(COMPAT,800)” on page 679  
The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. You should be aware of the implications of switching from OPMODE=(NEWFUNC,800) to OPMODE=(COMPAT,800)

## Preparing to migrate a single IBM MQ for z/OS queue manager

Follow the steps to prepare a single IBM MQ queue manager on z/OS for migration.

### About this task

To prepare to migrate an IBM MQ queue manager on z/OS, you need to carry out the detailed steps in this topic, using the links within this overview.

1. Make your existing queue manager ready for migration; see Step 1
2. Install the new code and make target libraries available to all MVS systems that are running queue managers, and grant access; see Step 2.
3. Perform a back up operation of each queue manager in your enterprise; see Step 3.
4. Review definitions of the user IDs for the queue manager(MSTR) and channel initiator (CHIN) address spaces; see Step 4
5. Restart your IBM MQ systems; see Step 4.
6. Review pageset zero usage before migration; see Step 5.
7. Migrate your Db2 tables, and repeat this step for each queue sharing group (QSG), if your enterprise uses QSGs; see Step 6
8. Add a new coupling facility (CF) structure definition and repeat this step for each QSG, if your enterprise uses QSGs; see Step 7.
9. Consider the migration of your application server; see Step 8
10. Configure Advanced Message Security (AMS); see Step 9

### Procedure

1. Make your IBM MQ configuration ready for migration.
  - a. Refer to the Preventive Service Planning (PSP) bucket for your version of IBM MQ ; see PSP Buckets - How to find them on Web.
  - b. Apply the migration and toleration PTFs to the version of the IBM MQ code that your enterprise uses; see IBM MQ Support, Migration PTFs.

Note that the “migration and toleration” PTFs are also known as the “backward migration and coexistence” PTFs; they are the same PTFs.

If you are unsure which migration PTFs you require, run the following command SMP/E:

```
REPORT MISSINGFIX ZONES(mqgtzone) FIXCAT(IBM.Coexistence.MQ.V8R0M0)
```

See FIXCAT and IBM MQ Migration Installation for further information.

**Attention:** If a PTF requires a rebind of Db2 plans, the PTF is shipped with ++HOLD(ACTION), indicating the need for this process. In such a case, see Migrating Db2 tables to bind the plans before starting migration.

Other FIXCAT categories are listed in IBM Fix Category Values and Descriptions.

There is an additional category of TargetSystem-RequiredService.MQ.V8R0M0 enabling other products to run with IBM MQ Version 8.0.

2. Install the new code and make target libraries available to all MVS systems that are running queue managers, and grant access. You must carry out the following procedure for each MVS system.

- a. Copy the IBM MQ target libraries to the system, and install the early code for the new version (once for each MVS system). Activate the code for all of the queue managers on each MVS system that is running queue managers.

This updates the LPA. See Update the z/OS link list and LPA for more information.

- b. APF authorize the load libraries and grant access to the datasets using your external security system. See APF authorize the IBM MQ load libraries for more information.
- c. Copy the file system zFS and mount it read only.

You only need zFS, or older HFS, if the IBM MQ for z/OS Unix System Services Component is installed. See the Program Directory for further information.

Refresh all the queue managers so that they use the new early code using the command REFRESH QMGR TYPE(EARLY). See REFRESH QMGR for more information.

3. Perform a back up operation for each queue manager in your enterprise, so that you have a before copy of all objects and JCL before you make any changes. This makes rolling back to the current system easier, if you require to do so.
  - a. Backup your IBM MQ defined objects, for example using CSQUTIL COMMAND MAKEDEF(..) See Issuing commands to Websphere MQ (COMMAND) for more information.
  - b. Backup:
    - The MSTR and CHINIT started procedure jobs
    - The Initialization input datasets used in the CSQINP1 and CSQINP2 concatenations
    - The system parameter module (ZPARM) libraries
    - Other tasks as necessary.

**Note:** You might also make a back up of page sets, BSDSs, and active logs as a fallback option. See How to backup and recover pagesets for more information on backing up IBM MQ resources.

4. Check that MSTR and CHIN address spaces run under user IDs that have OMVS segments defined, with a valid UID, to enable calling Unix System Services (USS).
5. Restart your IBM MQ system to run with the migration and toleration PTFs.
  - a. Restart the queue managers and monitor the whole system in your enterprise closely to ensure that there are no issues. Depending on the size and complexity of your enterprise this can take a considerable length of time, so you must plan for this in your migration schedule.
6. Review the usage of pageset 0. Note that you can ignore this step if your enterprise is already using IBM MQ V7.1. Issue the operator command /cpf DISPLAY USAGE PSID(0), where **cpf** is the command prefix for the queue manager's subsystem, to get a report on pageset 0 usage.

The size of queue definitions increased in IBM MQ V7.1. During migration to this version, or later versions of the product from an earlier version of the product, queue definitions stored on pageset 0 are rewritten.

The rewrite is carried out as a single transaction when the queue manager is first migrated to IBM MQ V7.1, or later.

Ensure that there is enough space available on pageset 0 to create a copy of the queue definitions while migration is taking place. Typically, 60% free space on pageset 0 before migration is sufficient. However, the use of EXPAND(SYSTEM) on the pageset definition allows for automatic expansion as required.

If there is insufficient space on pageset 0 during migration, the queue manager abends with completion code X'5C6' and reason code X'00C91900'.

7. Migrate your Db2 tables for each Db2 data sharing group.

You must do this for each Db2 data sharing group, as multiple QSGs can use the same Db2 tables.

You can use IBM provided samples shipped in the new version of the product to perform this task. Some Db2 table definitions are updated, and some new Db2 tables are created for the migrated version of the queue manager.

**Notes:**

- a. You must have applied the migration and toleration PTFs to all the queue managers, before migrating the Db2 tables.
- b. Every queue manager in the QSG needs to be restarted at the current release, with the PTFs applied.
- c. At no stage is an outage of the entire queue-sharing group required.
- d. Migrate your Db2 tables.

If the jobs described fail because of a Db2 locking problem, it might be due to contention for a Db2 resource. Locking is more likely, if the system is being heavily used. Resubmit the job later, preferably when the system is lightly used or quiesced.

See steps 5 and 6 of Set up the Db2 environment.

- e. Use the CSQ45\* jobs in the newest *thlqual.SCSQPROC* supplied with the version of the product to which you are migrating.

Note that the JCL to use depends on the highest version of IBM MQ in the Db2 tables.

**Attention:** If the Db2 tables have IBM MQ Version 8.0 queue managers, ignore the preceding steps, b and c.

- 1) If the Db2 tables have IBM WebSphere MQ Version 7.1 queue managers, use CSQ4571T. If the Db2 tables have IBM WebSphere MQ Version 7.0 queue managers, use CSQ4570T.

- 2) Customize the CSQ45\* sample.

The header information in CSQ45\* describes how to customize the sample.

- 3) Run the customized CSQ45\* job.

- 4) Customize the CSQ45BPL and CSQ45GEX samples, in *thlqual.SCSQPROC*

The header information in CSQ45BPL and CSQ45GEX describes how to customize the samples.

- 5) Run the customized jobs, CSQ45BPL and CSQ45GEX.

If you need to rebind the plans or packages, see Step 2 on page 567 in Migrating queue sharing groups from a previous release for further information.

- f. If you have multiple QSGs in the same data sharing group (DSG) you need to check each QSG to see if each member meets its migration criteria. Use sample JCL CSQ45MQS in conjunction with CSQ4571T.

See the JCL header description for further information.

8. Add the new coupling facility (CF) definition. Repeat this step for each QSG. Note that you can ignore this step if your enterprise is already using IBM MQ V7.1.

Starting with IBM MQ V701, a new CF structure is required; see Set up the coupling facility for information on how to add such a definition.

The correct process to migrate SYSTEM.QSG.CHANNEL.SYNCQ, from a normal application CF structure, to system CF structure CSQSYSAPPL structure is:

- a. Stop the channel initiator (CHINIT) on all QSG queue managers, so that no channels are running.
- b. Copy the messages in SYSTEM.QSG.CHANNEL.SYNCQ to a temporary dataset, using CSQUTIL COPY.
- c. Delete SYSTEM.QSG.CHANNEL.SYNCQ from the repository.
- d. Define SYSTEM.QSG.CHANNEL.SYNCQ with CFSTRUCT(CSQSYSAPPL). As this is a shared queue, it only needs to be defined once per QSG. Note that you can define this queue from any queue manager within the QSG.
- e. Reload the SYNCQ messages from the temporary dataset, back to the newly defined shared queue, using CSQUTIL LOAD.
- f. Perform the other migration steps, and then restart CHINIT to make the changes taking effect.

9. Migrate server applications.

Java or JMS applications running on the same host with IBM MQ connect to queue managers in bindings mode. This is a cross-memory connection. In this mode, applications need to update their STEPLIB concatenations, so that they can always load the highest version IBM MQ library in the system.

Note, that if a z/OS Java or JMS application is running under WebSphere Application Server, the application can use client mode as an alternative to bindings mode.

IBM MQ libraries include:

**thlqual.SCSQANLx**

This library contains error message information for your national language. The letter 'x' represents the letter for your national language.

**thlqual.SCSQAUTH**

This library contains the code that the applications use.

Server applications for IBM MQ can include:

- Batch applications
  - Control panels in ISPF
  - IMS
  - Interactive problem control system (IPCS)
  - RRS adapter, including Db2 stored procedures.
  - TSO
  - Additionally, WebSphere Application Server for z/OS, IBM Integration Bus, and CICS.
- a. You can use the "TSO ISRDDN ENQ ' thlqual.SCSQANLE'" command, replacing thlqual with the High Level Qualifier for your installation, to check which jobs are running with the specified library. You can then modify them accordingly.
  - b. Update STEPLIB in the application JCL, and refer to the new IBM MQ libraries.
  - c. Restart these applications. For further information, see:
    - Set up Batch, TSO, and RRS adapters
    - Setting up the IMS adapter
    - Set up the operations and control panels
    - Include the IBM MQ dump formatting member
  - d. Migrate other software, such as WebSphere Application Server, IBM Integration Bus, or CICS to use the version of IBM MQ that you need.
    - CICS

Update the IBM MQ libraries in the STEPLIB and DFHRPL concatenations of your CICS region JCL and restart CICS.

Up to, and including CICS 3.2, the connection between IBM MQ and CICS is provided by IBM MQ. You must change the SCSQCICS and SCSQAUTH libraries in the DFHRPL concatenation provided by IBM MQ.

After CICS 3.2, the connection between IBM MQ and CICS is provided by CICS libraries. Update the libraries, if you are using CICS Transaction Server for z/OS Version 3.2 or later. Without this change, you are not able to use the most recent IBM MQ features. You must change the SCSQCICS library in the DFHRPL concatenation provided by IBM MQ, and also the STEPLIB concatenation.

Create separate CICS started procedure JCL. For each CICS region that is connected to an IBM MQ queue manager, ensure that there is a separate CICS started procedure JCL.

This ensures that the modification of reference to a certain version of IBM MQ libraries in the CICS started procedure JCL only has impact for that single CICS region. In this way you can migrate one queue manager, and only the CICS region or regions connected to it, which makes staged migration possible.

CICS STEPLIB has thlqual.SCSQAUTH, and DFHRPL has thlqual.SCSQCICS, thlqual.SCSQLOAD, and thlqual.SCSQAUTH. For more information, see [Setting up the CICS- IBM MQ adapter](#)

- WAS for z/OS

If you are running in an application server environment where a bindings connection is being used, you need to update the WAS STEPLIB with IBM MQ libraries.

See [IBM MQ libraries and the WebSphere Application Server for z/OS STEPLIB](#) for further information.

You also need to configure the IBM MQ messaging provider with native libraries from the new version of the IBM MQ installation; see [Configuring the IBM MQ messaging provider with native libraries](#) for further information.

Use the latest level of native libraries in USS.

Note that you can make use of a DFP ALIAS for convenience. Create data set aliases such as MQM.SCSLOAD, and reference them in JCL. Map the aliases to the real data sets, such as MQM.V700.SCSLOAD or MQM.V710.SCSLOAD.

Change the aliases to switch between the two sets of target libraries. With the aliases, you can start applications or the queue manager when moving to a new release of IBM MQ without changing the STEPLIB JCL.

## 10. Advanced Message Security (AMS)

If the queue manager is configured to use Advanced Message Security (AMS) perform the steps in the [Preparing to migrate Advanced Message Security](#) section of the [Migrating Advanced Message Security](#) topic.


## Results

You have prepared your IBM MQ queue manager on z/OS for migration.

## What to do next

Follow the instructions in “[Migrating a single IBM MQ z/OS queue manager to the next release of the product](#)” to migrate the queue manager.

### Related information:

 [Program Directory for IBM MQ for z/OS](#)

## Migrating a single IBM MQ z/OS queue manager to the next release of the product

Carry out the instructions in this topic to migrate a single IBM MQ queue manager on z/OS,

### About this task

To migrate an IBM MQ queue manager on z/OS to a different release, you need to carry out the:

- Process described in “[Preparing to migrate a single IBM MQ for z/OS queue manager](#)” on page 554
- Detailed steps in this topic, using the links within this overview.
  1. Stop or disconnect the applications; see step 1 on page 559
  2. Stop the queue manager and its channel initiator; see step 2 on page 559
  3. Update STEPLIB for MSTR and the channel initiator; see step 3 on page 559
  4. Update the initialization input data sets; see step 4 on page 559
  5. Update the target version system parameter module (ZPARM); see step 5 on page 559
  6. Configure Advanced Message Security; see step 6 on page 559
  7. Review the security control of your system; see step 7 on page 559
  8. Start the queue manager; see Step 8 on page 559

9. Optionally, revert the queue manager to a previous version; see step 9 on page 560

## Procedure

1. Stop or disconnect all the applications using the queue manager (for example, CICS, IMS, or batch) and the IBM MQ channels that are connected to other queue managers.
2. Stop the queue manager and its channel initiator.
3. Update STEPLIB for MSTR and the channel initiator (CHINIT). Update the start procedure and CHINIT JCL.
  - a. Update your procedure to start the queue manager.

Change the STEPLIB for the queue manager to reference the new version of the libraries.  
See Create procedures for the IBM MQ queue manager.

IBM MQ now uses z/OS memory objects above the bar for some functions. You must allow the queue manager to access storage above the bar.

Your installation might have customized the SMFPRMxx member of SYS1.PARMLIB, or the **IEFUSI** exit to provide a default limit for jobs using virtual storage above the 2 GB bar. Check these limits give sufficient memory for a queue manager. A reasonable starting allocation is 2 GB. The message CSQY220I displays the amount of virtual storage currently used and available.

If your installation does not have a default limit for storage above the bar, or if you want to use a different limit for your queue manager, you can provide a queue manager-specific restriction on the amount of virtual storage available above the bar for memory objects by coding a **MEMLIMIT** parameter on the JCL of the queue manager stored procedure, xxxxMSTR, for example:

```
//PROCSTEP EXEC PGM=CSQYASCP,REGION=0M,MEMLIMIT=2G
```

MEMLIMIT defines memory available above the bar; see Address space storage
  - b. Update your procedures for the channel initiator.

Change the STEPLIB for the channel initiator to reference the new level of the product libraries.  
See Create procedures for the channel initiator.
4. You must allow the queue manager to access storage above the bar because IBM MQ uses memory above the bar. If insufficient storage is available above the bar, the queue manager reports this when starting, and stops.

See step 3a for further information on **MEMLIMIT**.
5. Optionally, update the target-version system parameter module (ZPARM). Tailor the target version (ZPARM) sample to use new MQ libraries and generate new ZPARM.

Some new parameters might be added in the target version; the current version of the IBM MQ libraries cannot recognize them.

Configure the OPMODE parameter, for example OPMODE=(COMPAT,800) to allow the queue manager to revert if necessary. For further details, see OPMODE.

ZPARM is forward compatible, as it simplifies the forward migration process. This means that any attributes introduced in the new release will have their default values.

This step is optional during queue manager migration; you can complete this task after you confirm that no problems occurred after migration. This option gives you an easier approach when backward migration is needed.
6. If the queue manager is configured to use Advanced Message Security (AMS), perform the steps in Migrating Advanced Message Security.
7. Review your security control for queue sharing groups, the channel initiator, and all queue managers accessing the coupling facility list structures.
8. Start the queue manager. Test that everything is working correctly and, if it is, start the channel initiator. If there is a problem starting the queue manager, consider reverting the queue manager to a previous version; see step 8.

9. If a problem occurs when starting the queue manager, you might need to consider backward migration; see Reverting a queue manager to a previous release.

## Results

You have migrated your IBM MQ for z/OS queue manager to the latest release.

## What to do next

Follow the instructions in “Post migration tasks” on page 564 to complete the migration process.

### **z/OS: CSQINP1 and CSQINP2 input data sets changed**

The CSQINP1 and CSQINP2 initialization input data sets changed in Version 7.1. The data sets include more samples and the contents of some samples have been moved to other samples. Particular changes to take note of are the commands to define queues to hold publish/subscribe state information. The commands must be in the right order.

### **Important changes to initialized input data sets since IBM WebSphere MQ Version 7.1**

#### **CSQ4INSM**

Added for IBM MQ Advanced Message Security support

#### **CSQ4INSG**

Add one **AUTHINFO** object SYSTEM.DEFAULT.AUTHINFO.IDPWOS for connection authentication support.

Some channel objects and topic objects are modified with new attributes, for example, **STATCHL** and **CLROUTE**.

#### **CSQ4INST**

The default system subscription, SYSTEM.DEFAULT.SUB, moved from CSQ4INSG to CSQ4INST in IBM WebSphere MQ Version 7.1.

#### **CSQ4INSX**

Add one model queue SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE for multiple cluster transmission queue support.

Review the changes, and update the customized versions you are currently using as required, if your enterprise is using IBM WebSphere MQ Version 7.0.

The queue manager uses queues to hold persistent state information about publish/subscribe. Durable subscriptions are held as messages on SYSTEM.DURABLE.SUBSCRIBER.QUEUE and retained publications on SYSTEM.RETAINED.PUB.QUEUE.

The ordering of the definitions of these queues is important. Member CSQ4MSTR of **SCSQPROC** shows the necessary ordering of the supplied definitions in the CSQINP2 concatenation. The default system subscription, SYSTEM.DEFAULT.SUB, requires SYSTEM.DURABLE.SUBSCRIBER.QUEUE that in turn requires the storage class SYSLNGLV, which is defined in CSQ4INYS. If you migrating from a previous release, and modifying customized procedures, define these resources in the following order:

1. Storage class SYSLNGLV. It requires a defined mapping to a defined page set.
2. SYSTEM.DURABLE.SUBSCRIBER.QUEUE
3. SYSTEM.DEFAULT.SUB

#### **Note:**

The changes to CSQINP1 and CSQINP2 are required in Create procedures for the IBM MQ queue manager and Customize the initialization input data sets.



## **z/OS Migrating a queue manager to mixed case security**

Follow these steps to migrate a queue manager to mixed case security. You review the level of security product you are using and activate the new IBM MQ external security monitor classes. Run the **REFRESH SECURITY** command to activate the mixed-case profiles.

### **Before you begin**

1. Install a level of the security product that supports mixed case security.
2. Apply any updates required by IBM MQ.
3. Install and activate the new IBM MQ external security monitor classes.

### **About this task**

Follow these steps to convert a queue manager to mixed case security.

### **Procedure**

1. Copy all your existing profiles and access levels from the uppercase classes to the equivalent mixed case external security monitor class.
  - a. MQADMIN to MXADMIN.
  - b. MQPROC to MXPROC.
  - c. MQNLIST to MXNLIST.
  - d. MQQUEUE to MXQUEUE.
2. Start the queue manager.

The queue manager SCYCASE attribute is set to UPPER.
3. Change the value of the SCYCASE attribute to MIXED.

```
ALTER QMGR SCYCASE(MIXED)
```
4. Activate your existing security profiles.

```
REFRESH SECURITY(*) TYPE(CLASSES)
```
5. Test that your security profiles are working correctly.

### **What to do next**

Review your object definitions and create new mixed case profiles as appropriate, using **REFRESH SECURITY** as required to activate the profiles.

## **z/OS: Migrating IBM MQ Advanced Message Security**

IBM MQ Advanced Message Security for z/OS (AMS) is a separately licensed enabling product that extends IBM MQ to provide a high level of protection for sensitive data flowing through the IBM MQ network using a public key cryptography model.

In IBM MQ for z/OS releases prior to Version 8.0, AMS was provided as a separate product. This topic describes the tasks required to migrate the AMS configuration on z/OS from that used in version 7 and earlier, to that used in Version 8.0. These steps supplement those required to migrate a single IBM MQ for z/OS queue manager where AMS is not configured. AMS must be migrated at the same time as the queue manager, it is not supported to use IBM MQ Advanced Message Security Version 7.0.1 with IBM MQ for z/OS Version 8.0.

To enable AMS on a newly created IBM MQ for z/OS queue manager, or on a queue manager that has already been migrated to Version 8.0, see IBM MQ Advanced Message Security for z/OS.

## Preparing to migrate Advanced Message Security

To prepare to migrate an IBM MQ queue manager on z/OS using AMS Version 7.0.1 or earlier, you must perform the steps in this section in addition to those listed in “Preparing to migrate a single IBM MQ for z/OS queue manager” on page 554.

1. Install the IBM MQ Advanced Message Security for z/OS enabling product and make the target libraries available to all MVS systems that are running queue managers that will use AMS. You must carry out the following procedure for each MVS system:
  - a. Copy the AMS target libraries to the system.
  - b. APF authorize the thlqual.SDRQAUTH target library and grant access to this data set using your external security system, see Task 2: APF authorize the IBM MQ load libraries.
  - c. Ensure the LPA contains the AMS module CSQ0DRTM, see Task 3: Update the z/OS link list and LPA.
  - d. Ensure the program properties table (PPT) contains an entry for CSQ0DSRV, see Task 4: Update the z/OS program properties table.
2. For each queue manager, set up the started task user for the AMS address space. In AMS Version 7.0.1 two address spaces are used, one for the main task and another for the data services task. In Version 8.0 these are combined in to a single address space called qmgrAMSM. Either set up a new user for the Version 8.0 AMS address space, or grant additional authorities to one of the existing AMS started task users. See Task 25: Set up the started task user Advanced Message Security for information on how to set up the started task user. If you do not use the existing data services address space user you will need to replicate the **drq.ams.keyring** key ring for the user ID associated with the Version 8.0 qmgrAMSM address space. See Using certificates on z/OS for information on how to set up the AMS key ring.

## Migrating Advanced Message Security

To migrate an IBM MQ queue manager on z/OS using AMS Version 7.0.1 or earlier, before restarting the queue manager you must perform the steps in this section in addition to those listed in “Migrating a single IBM MQ z/OS queue manager to the next release of the product” on page 558.

1. Take a copy of the qmgrAMSM task for IBM MQ Advanced Message Security (AMS) Version 7.0.1, in case you need to revert to your previous system.  
See “Backward migration of Advanced Message Security” on page 563 for more information.
2. Configure the queue manager to use AMS by updating the system parameter module to set SPLCAP(YES) using CSQ6SYSP, see Task 17: Tailor your system parameter module and Using CSQ6SYSP.
3. Create or update the started task procedure for the qmgrAMSM address space, see Task 24: Create procedures for Advanced Message Security.

## Post migration tasks for Advanced Message Security

After you have migrated an IBM MQ queue manager on z/OS that uses AMS you must perform the following tasks.

1. In Version 8.0, the AMS address space is started and stopped automatically by the queue manager. If you have automation to manage the main task and data services task for AMS Version 7.0.1, or earlier, this should be removed. You must also review any automated console commands for AMS because some have changed in version 8.
2. Delete the started task procedures for the Version 7.0.1 data services task and the Version 7.0 main task if these were not called qmgrAMSM.

## Backward migration of Advanced Message Security

If you are an AMS user, and you backward migrate your queue manager from Version 8.0 to a version 7 release, additional actions are required to revert AMS to version 7.

### Considerations when migrating

You should ensure that your previous setup is in place and that tasks Updating the z/OS LPA to Updating your system DIAG member have been carried out.

Ensure that the user ID associated with the version 7 data-services address spaces has access to `drq.ams.keyring`, and that `drq.ams.keyring` has the same connected certificates as the Version 8.0 `qmgrAMSM` user ID.

### Performing the migration

When you have completed the previous tasks, you can migrate your queue manager backwards in the normal way.

Manually start, or reintroduce automation for starting, the AMS main and data services address spaces.

See Starting IBM MQ Advanced Message Security for further information.

### z/OS: Reverting a queue manager to a previous release

After migrating to IBM MQ for z/OS Version 8.0, from either Version 7.0.1 or Version 7.1.0, you can backward migrate, or fallback, to the version you were using prior to migration. Backward migration Program Temporary Fixes (PTFs) are available for both Version 7.0.1 and Version 7.1.0.

### Before you begin

In general, after fallback to IBM WebSphere MQ Version 7.0, new attributes of IBM MQ objects introduced at Version 8.0 will be removed. The APAR that supplies these PTFs, documents specific information relating to fallback to IBM WebSphere MQ Version 7.0.1 or to IBM WebSphere MQ Version 7.1.0.

Switching back to running a queue manager with the target libraries of a previous version is possible if **DISPLAY** **SYSTEM** returns `COMPAT=vrm` where `vrm` is the level of the previous version. If it does, you can go back to using your customization and startup procedure for the queue manager from that version

- The queue manager compatibility level must be `7rm`. If the queue manager has never been run with **OPMODE** set to `(NEWFUNC,800)`, the compatibility level is `7rm`.
  - **DISPLAY** **SYSTEM** returns `OPMODE COMPAT,7rm`.

where

**7** is the version number `v` of the product.

**r** is the release number of the product.

**m** is the modification number of the product.

- Before migrating your queue manager to the latest version, with that version of target libraries, you applied all the migration and toleration PTFs to the queue manager on your previous version. The queue manager then started successfully with those PTFs at that previous version. This is a requirement before you can revert your queue manager to the original version.
- You saved the queue manager customization macros and JCL for running with the Version 7.0 target libraries.

You can recreate the customization for version `7.r`, if the originals are not available to you.

## About this task

To restart the queue manager, so that it runs at the version where it was migrated from, just requires that you switch back to using the libraries for the previous version.

Note that it is not necessary to rollback the early code for this installation when reverting your queue manager to an earlier version.

## Procedure

1. Stop the listener, channel initiator, and queue manager.
2. Switch back to use the MSTR and CHINIT started procedure JCLs with version 7.r libraries.  
In case data set aliases are used for load libraries, switch the aliases to refer to version 7.r libraries. For example, an alias named MQM.MQP1.SCSLOAD, referring to MQM.MQV800.SCSLOAD, needs to change to refer to MQM.MQV7xx.SCSLOAD.
3. Restart queue manager, using the system parameter module (CSQZPARM) used with IBM MQ version 7.r prior to migration, and linking to the version 7.r code.  
Until you have verified the startup, start the queue manager, channel initiator, and listener separately, checking for errors on the console after each component is started. If the startup runs cleanly, combine the startup of all three components in the production environment.
  - a. Start the queue manager.
  - b. Start the channel initiator.
  - c. Start the listener.
4. Verify correct functioning of existing applications.

## Results

If the queue manager cannot be reverted to the previous release by following the preceding procedure, for example, because it has been started with OPMODE set to (NEWFUNC,800), the queue manager can only be reverted to the previous release by recovering the page sets, BSDSs, and active logs from back up copies taken before the migration to IBM MQ for z/OS Version 8.0.

All updates made since the back up was taken will be lost. See How to backup and recover pagesets for more information on backing up IBM MQ resources.

## Post migration tasks

Follow the steps to perform the tasks you need to carry out after migrating a single IBM MQ queue manager on z/OS,

## About this task

After you have migrated an IBM MQ queue manager on z/OS you need to carry out the detailed steps in this topic, using the links within this overview.

1. Check the changes in behavior made by default configuration changes; see Step 1 on page 565
2. Modify the backup jobs to refer to the target version of IBM MQ libraries; see Step 2 on page 565
3. Configure Advanced Message Security; see 3 on page 565
4. Perform a full regression test; see Step 4 on page 565
5. Update the ZPARM module if you have not already done so; see Step 5 on page 565
6. Set OPMODE to NEWFUNC; see Step 6 on page 565
7. Exploit the new functions provided by the migrated queue manager; see Step 7 on page 565
8. Consider client application migration; see Step 8 on page 565

## Procedure

1. Check the changes in behavior made by default configuration changes. The default values of some properties might have been changed in the new version, which can lead to changes in behavior. SHARECNV allows multiple connections to the queue manager to permit the use of the same TCP/IP connection. If a client is using Version 6 code to connect to a version 7, or later, IBM MQ queue manager, SHARECNV is set to 0 automatically; see Default behavior for more details about this change.  
On z/OS, you can reverse queue manager migration as long as you have not enabled new function. You enable new function by setting the **OPMODE** parameter to (NEWFUNC,800) ; see OPMODE for more information.
2. Modify backup, and other administrative, jobs to refer to the target version of IBM MQ libraries, such as backup IBM MQ objects and MAKEDEF jobs. For example using CSQUTIL COMMAND MAKEDEF(.); see Using the COMMAND function of CSQUTIL.  
You should also backup channel authentication records, which were introduced in IBM WebSphere MQ Version 7.1.0.
3. If the queue manager is configured to use Advanced Message Security (AMS) perform the steps in the Post migration tasks for Advanced Message Security section of the Migrating Advanced Message Security topic.
4. Perform a full regression test.
5. Update the ZPARM module if you have not already done so. See Update the ZPARM module for further information.
6. Set OPMODE in ZPARM JCL to NEWFUNC and recompile the JCL. For more information about NEWFUNC, see OPMODE.
7. Exploit the new functions provided by the migrated queue manager. Your queue manager has been fully migrated to a new version level, and you can take benefit of new capability now.  
Review What's new in IBM MQ Version 8.0 and check which features best serve your business needs. Plan your action to develop new applications, or changing configurations, to enable those features.
8. Migrate client applications. Client applications running on z/OS, or other distributed platforms, can be considered any time throughout the migration phase.  
The client libraries need to be at the same level as the lowest IBM MQ queue manager level to which they connect. So, if a client can connect to Version 701, Version 710, or Version 800 IBM MQ queue managers, the client needs to be at Version 701. Once all the IBM MQ queue managers have been migrated to Version 800 you can migrate the clients to Version 800.  
For further information, see "Migrating an IBM MQ MQI client on UNIX platforms, and Windows to the latest version" on page 475 and select the platform that you require.

## Results

You have completed the migration of a single IBM MQ for z/OS queue manager.

---

## z/OS: Adding a new queue-sharing group to an existing Db2 data sharing group in the latest version

Follow these steps to add a new queue-sharing group to an existing Db2 data sharing group in the latest version of the product. You must apply the migration and toleration PTFs to queue managers, in the previous version, in any of the queue-sharing groups before adding a queue-sharing group.

### Before you begin

1. Review your Db2 data-sharing requirements. A single Db2 data-sharing group can be used to support multiple IBM MQ queue-sharing groups.
2. You can add a new queue-sharing group to a Db2 data-sharing group that already supports IBM MQ queue-sharing groups containing queue managers for the previous version. You must ensure that the migration and toleration PTFs have been applied. The Db2 tables used by IBM MQ must be configured for the latest version queue managers.

### About this task

Queue-sharing group migration affects steps 4 on page 491, 8 on page 492, 11 on page 492, and 12 on page 493 of “z/OS: Review and modify queue manager customizations from the previous release” on page 489

### Procedure

1. Customize the CSQ4570T and CSQ4571T samples, in *thlqual*.SCSQPROC, supplied with the latest version of the IBM MQ for z/OS product.
  - The header information in CSQ4570T and CSQ4571T describes how to customize the samples.
  - Delete or bypass the step that runs MIGRATE QSG.
2. Run the customized CSQ4570T and CSQ4571T jobs.
3. Set up the coupling facility.
  - See Task 10: Set up the coupling facility.
4. Customize and include the initialization input sample *thlqual*.SCSQPROC(CSQ4INSS) in the CSQINP2 data set.
  - See step 11 on page 492 in “z/OS: Review and modify queue manager customizations from the previous release” on page 489.
5. Add the IBM MQ entries to the Db2 data-sharing group using the **CSQ5PQSG** program.
  - See Task 16: Add the IBM MQ entries to the Db2 data-sharing group.
6. Tailor the system parameter module to add Db2 data-sharing group and IBM MQ queue-sharing group information.
  - See step 12 on page 493 in “z/OS: Review and modify queue manager customizations from the previous release” on page 489.

---

## z/OS: Migrating queue sharing groups from a previous version of the product



You can migrate one or more existing queue-sharing groups containing queue managers in a previous version of a product to the latest version. At no stage is an outage of the entire queue-sharing group required.

### Before you begin

1. Read the topic, “Queue-sharing group migration” on page 480, and its related references, especially “z/OS: Queue sharing group coexistence” on page 583.

### About this task

Migrating each queue manager comprises the bulk of the work of migrating a queue-sharing group. Approach migrating a queue sharing group as requiring some extra tasks that must be performed during the migration of each queue manager. A good approach is to create a migration plan incorporating queue-sharing group migration; see “z/OS: Migration planning to the latest release” on page 486.

Queue-sharing group migration affects steps 4 on page 491, 8 on page 492, 11 on page 492, and 12 on page 493 of “z/OS: Review and modify queue manager customizations from the previous release” on page 489

### Procedure

1. Ensure that all members of the queue-sharing group have been started at the same version, before migrating any queue managers in the queue-sharing group to the latest version.

Do this, by migrating the queue managers at the oldest version in the queue-sharing group to the same version as the other queue managers. For example, if the queue-sharing group currently contains queue managers at Version 7.0.1 and Version 7.1, migrate the Version 7.0.1 queue managers to Version 7.1, before migrating any queue managers in the queue-sharing group to Version 8.0.

2. Apply IBM MQ for z/OS migration and toleration<sup>13</sup> PTFs for the latest version of the product to the earlier version code; see IBM MQ Support, Migration PTFs.

- a. Apply the PTFs to the libraries of the earlier version of the product. If you need to rebind a plan or package, follow the instructions in the PTFs, and rebind the packages using the instructions in CSQ45BPK.

If you rerun the bind of the plan, or package, while queue managers are active, the bind fails with a locking problem if any queue manager in the DSG using the plan is active.

You can either shut down the queue managers using the plans, or suspend the queue managers use of Db2. See Suspending a connection to Db2 for further information.

- b. Perform the additional hold action tasks of binding new and changed DBRMs into plans
- c. Stop and restart each queue manager so that it picks up the new code level.
- d. Perform testing of the new code level.

These steps can be performed at any time in preparation for a migration to the latest version of IBM MQ for z/OS, or as part of normal maintenance. It is not dependent on the latest version being available.

Migrating an earlier version queue manager to a later version queue manager within a queue sharing group is restricted. The restrictions are, that all the earlier version queue managers in the queue-sharing group must have been started at the *same* earlier version, and that all the earlier

---

13. The “migration and toleration” PTFs are also known as the “backward migration and coexistence” PTFs. They are the same PTFs.

version queue managers have had the “earlier version backwards migration from the latest version, and coexistence with the latest version” PTFs applied.

After a queue manager for the latest version has been started in a queue-sharing group, starting an earlier version queue manager is restricted. You cannot start an earlier version queue manager as a member of the group, unless it has migration and toleration PTFs applied.

The latest version requires new Db2 tables, and additional changes to existing Db2 tables. The PTF changes some of the Db2 operations performed by the earlier version queue manager. The changes make an earlier version queue manager compatible with the latest version.

The PTF contains a new set of Database Request Modules (DBRM). After binding Db2 with these DBRMs, you have two sets of plans: one set for queue managers without the PTFs and the other set for queue managers with the PTFs applied.

### 3. Migrate your Db2 tables.

You must have applied the migration and toleration PTFs to all the queue managers in the queue-sharing group before migrating Db2 tables.

If the jobs described fail because of a Db2 locking problem, it might be due to contention for a Db2 resource. Locking is more likely, if the system is being heavily used. Resubmit the job later, preferably when the system is lightly used or quiesced.

You can either migrate the Db2 tables one queue-sharing groups at a time, or all queue-sharing groups at the same time. For more information, read the header information in the jobs.

- Migrate the tables for all the queue-sharing groups at the same time.
  - a. Customize the CSQ4570T and CSQ4571T samples, in *thlqual*.SCSQPROC, supplied with the latest version of the IBM MQ for z/OS product.

The header information in CSQ4570T and CSQ4571T describes how to customize the samples.
  - b. Run the customized CSQ4570T and CSQ4571T jobs.
  - c. Customize the CSQ45BPL and CSQ45GEX samples, in *thlqual*.SCSQPROC

The header information in CSQ45BPL and CSQ45GEX describes how to customize the samples.
  - d. Run the customized jobs, CSQ45BPL and CSQ45GEX.

The customized jobs are run as part of step 8 on page 492 of “z/OS: Review and modify queue manager customizations from the previous release” on page 489.

This step binds the latest version DBRMs into plans, and grants execute authority to them.
- Migrate the tables for one queue-sharing group at a time.
  - a. Customize the CSQ4570T and CSQ4571T samples, in *thlqual*.SCSQPROC, supplied with the latest version of the IBM MQ for z/OS product.

The header information in CSQ4570T and CSQ4571T describes how to customize the samples.

Edit step 8 on page 492 of “z/OS: Review and modify queue manager customizations from the previous release” on page 489 that executes the MIGRATE QSG function, specifying the name of the first queue-sharing group that is to be migrated.
  - b. Run the customized CSQ4570T and CSQ4571T jobs.
  - c. Customize the CSQ45BPL and CSQ45GEX samples, in *thlqual*.SCSQPROC

The header information in CSQ45BPL and CSQ45GEX describes how to customize the samples.
  - d. Run the customized jobs, CSQ45BPL and CSQ45GEX.

The customized jobs are run as part of step 8 on page 492 of “z/OS: Review and modify queue manager customizations from the previous release” on page 489.

This step binds the latest version DBRMs into plans, and grants execute authority to them.



---

## Migrating a queue manager cluster

Migrate a queue manager cluster by migrating each of the queue managers in the cluster. No changes since Version 6.0 specifically effect the migration of queue manager clusters. But you must consider what the effect is of migrating some queue managers in a cluster, before all the queue managers are migrated.

### Before you begin

Check that no cluster-specific migration issues are identified for the migration you are intending to perform. No cluster-specific changes have been made that affect migration between Version 7.0.1 and later versions.

### Procedure

Consider the following issues that relate to migrating a queue manager cluster:

- Minimizing application outages.
- Measuring and verifying migration success and planning for backward migration if there are any migration problems.
- Taking advantage of new IBM MQ features.
- Managing the migration of a cluster in the context of the wider IBM MQ network and the systems architecture of your organization.

### Steps

Some techniques for migrating a cluster with minimal planned and unplanned outages are described in the following topics:

#### Related concepts:

“How mixed version cluster repositories are updated” on page 479

Repositories store records for an object in a cluster in the version of the record format that matches the version of the queue manager hosting the repository. Repository queue managers forward object records, before they are stored, in the format that they are received in. The recipient ignores fields from a newer version, and uses default values for fields that are not present in the record.

“Queue manager cluster migration” on page 478

You can migrate queue managers in a cluster all at once, or one at a time, which is called a staged migration. Migrate full repository queue managers in a cluster before partial repository queue managers.

## Migrating a queue manager cluster: Create a plan

Before carrying out the migration of a queue manager cluster, plan what you are going to do. Identify the roles that different queue managers play in the cluster, and decide in what order to migrate the queue managers.

### Procedure

- What queue manager and application migration issues must be dealt with between the old and new versions?
- What system architecture and change control procedures must you consider?
- Consider migration questions specific to clusters, such as migrating full repositories first, and migrating overlapping clusters.
- Are any of the queue managers in a queue-sharing group, or part of a high-availability solution?
- Is the cluster a publish/subscribe cluster? Which queue manager is a cluster topic host?
- Decide whether to carry out a staged migration, or migrate all queue managers at the same time.
- Do you have a test system to migrate, and a production system?
- Document and test the plan before migrating production queue managers.

### Related concepts:

“Application migration and interoperation” on page 477

IBM MQ supports running applications compiled and linked against previous versions of IBM MQ, with later levels of IBM MQ.

“Migrating a queue manager in a high-availability configuration” on page 481

Follow standard procedures to migrate a queue manager that is part of a high-availability configuration. (On platforms other than z/OS.)

“How mixed version cluster repositories are updated” on page 479

Repositories store records for an object in a cluster in the version of the record format that matches the version of the queue manager hosting the repository. Repository queue managers forward object records, before they are stored, in the format that they are received in. The recipient ignores fields from a newer version, and uses default values for fields that are not present in the record.

“Queue manager cluster migration” on page 478

You can migrate queue managers in a cluster all at once, or one at a time, which is called a staged migration. Migrate full repository queue managers in a cluster before partial repository queue managers.

“Queue manager migration” on page 472

After upgrading an installation, queue manager migration might be required. Migration takes place when you start a queue manager.

 “Queue-sharing group migration” on page 480

You can combine queue managers from different releases in a queue-sharing group. Limit the time you manage a mixed group to only as long as it takes to migrate all the queue managers to the same command level. You cannot combine a queue manager at Version 8.0 or later in the same queue-sharing group as queue managers earlier than Version 7.0.1. You must update all queue managers in a queue-sharing group with a coexistence PTF, before migrating any of them. All queue managers in the queue-sharing group must be at the same version before any queue manager can be migrated to Version 8.0.

### Related information:


Availability of cluster topic host queue managers

## Migrating a queue manager cluster: Create a backout plan

Before performing a migration, decide on a backout plan in case of failure.

### Before you begin

What backout capabilities do the queue managers in the cluster support?

 If the libraries of the earlier level of IBM MQ include the appropriate PTFs to be able to backward migrate, and NEWFUNC mode is not enabled at the higher level, queue managers running on z/OS can be reverted to an earlier level by changing the load libraries.

On other platforms, the only backout option is to restore a queue manager to a previous state. In restoring a queue manager, you lose any persistent changes since the queue manager started running at the new level.

### About this task

The backout plan must consider how to maintain the availability of the cluster. It must deal with any issues arising from migrating a queue manager in the cluster.

### Procedure

The backout plan must describe the following points:

- What constitutes a successful migration.

- The conditions that trigger the backout procedure.
- Alternative backout actions, such as:
  1. Suspending a queue manager from the cluster.
  2. Backward migration
  3. Keeping a queue manager offline until an external issue is resolved.

#### Related concepts:

“Reverting a queue manager to a previous version” on page 473

You can remove an upgrade before you have started a queue manager. After a queue manager has been started, if you remove the upgrade, the queue manager will not work.

## Migrating a queue manager cluster: Migrating one cluster queue manager

Follow these steps to migrate a single queue manager in a cluster. Base your cluster migration plan on applying these steps to each queue manager in the cluster.

### Procedure

1. Suspend the queue manager that you want to migrate from the cluster:
  - a. Issue the **MQSC** command:
 

```
SUSPEND QMGR CLUSTER(cluster name)
```
  - b. Check that no messages are sent to the queue manager.
 

You must close any application that continues to send messages to this queue manager. The cluster workload algorithm might choose the suspended queue manager. If there are no other valid destinations, or if an application has an affinity with the queue manager, it might select the queue manager.
2. Save a record of all cluster objects known by this queue manager. This data is used after migration to check that objects have been migrated successfully.
  - a. Issue the command to view cluster queue managers.
 

```
DISPLAY CLUSQMGR(*)
```
  - b. Issue the command to view cluster queues.
 

```
DISPLAY QC(*)
```
  - c. Issue the command to view cluster topics.
 

```
DISPLAY TCLUSTER(*)
```
3. Save a record from the full repository of its view of the cluster objects owned by this queue manager. The record is used after migration to check that objects have been migrated successfully.
  - a. Issue the command on the full repositories to display this queue manager.
 

```
DISPLAY CLUSQMGR(migrated queue manager name)
```
  - b. Issue the command on the full repositories to display the cluster queues for this queue manager
 

```
DISPLAY QC(*) WHERE(CLUSQMGR EQ migrated queue manager name)
```
  - c. Issue the command on the full repositories to display the cluster topics for this queue manager.
 

```
DISPLAY TCLUSTER(*) WHERE(CLUSQMGR EQ migrated queue manager name)
```
4. Migrate the queue manager.
 

Do one of the queue manager migration tasks, depending on the platform; see “Migrating a queue manager to the latest release” on page 512.

The queue manager migration process is, in outline:

  - a. Stop the queue manager.
  - b. Take a backup of the queue manager.
  - c. Install the new version of IBM MQ.
  - d. Restart the queue manager.

5. Ensure that all cluster objects have been migrated successfully.
  - a. Issue the command to view cluster queue managers and check the output against the data saved before migration.  
DISPLAY CLUSQMGR(\*)
  - b. Issue the command to view cluster queues and check the output against the data saved before migration.  
DISPLAY QC(\*)
  - c. Issue the command to view cluster topics and check the output against the data saved before migration.  
DISPLAY TCLUSTER(\*)
6. Check that the queue manager is communicating with the full repositories correctly.
7. Check that cluster channels to full repositories can start.
8. Check that the full repositories still have information about the migrated cluster queue manager, its cluster queues, and its cluster topics.
  - a. Issue the command on the full repositories and check the output against the data saved before migration.  
DISPLAY CLUSQMGR(*migrated\_queue\_manager\_name*)
  - b. Issue the command on the full repositories and check the output against the data saved before migration.  
DISPLAY QC(\*) WHERE(CLUSQMGR EQ *migrated\_queue\_manager\_name*)
  - c. Issue the command on the full repositories and check the output against the data saved before migration.  
DISPLAY TCLUSTER(\*) WHERE(CLUSQMGR EQ *migrated\_queue\_manager\_name*)
9. Test that applications on other queue managers can put messages to queues owned by the migrated cluster queue manager.
10. Test that applications on the migrated queue manager can put messages to the queues owned by other cluster queue managers.
11. Resume the queue manager by issuing the following command:  
RESUME QMGR CLUSTER(*cluster name*)
12. Closely monitor the queue manager and applications in the cluster for a while.

**Related concepts:**

“Queue manager migration” on page 472

After upgrading an installation, queue manager migration might be required. Migration takes place when you start a queue manager.

**Related information:**

DISPLAY CLUSQMGR

DISPLAY QUEUE

RESUME QMGR

SUSPEND QMGR

## Migrating a queue manager cluster: Migrating the test system

Migrate each queue manager in the test system.

### About this task

For each queue manager in the test system, in the order defined in the migration plan you developed in “Migrating a queue manager cluster: Create a plan” on page 569, migrate, and test the queue manager.

## Migrating a queue manager cluster: Migrating the production system

Migrate each queue manager in the production system.

### About this task

For each queue manager in the production system, in the order defined in the migration plan you developed in “Migrating a queue manager cluster: Create a plan” on page 569, migrate, and test the queue manager.

---

## Windows: Migrating an MSCS configuration

Migrate queue managers in MSCS configuration one node at a time, following these instructions.

### About this task

These steps are required for a rolling upgrade with a minimum amount of downtime. You must always upgrade an offline node with no online IBM MQ resources. In an Active/Passive configuration, if the node is Passive, you must ensure it cannot be switched to Active during the upgrade process.

The example, “Migrating a four-node MSCS cluster from an earlier version of the product to the latest version,” shows this procedure applied to a four-node cluster.

### Procedure

1. Modify the possible owners of the IBM MQ resource to encompass only the Active node or nodes. With no owners assigned to Passive nodes, the IBM MQ resource that is being migrated cannot be activated.
2. Ensure that the group containing the IBM MQ resource is currently on one of the nodes defined as a possible owner. The group must include any applications connecting to the queue manager resource.
3. Stop the cluster service on the node being migrated. The MSCS cache is cleared of any IBM MQ DLLs that have been registered.
4. Migrate the selected node by following the standard instructions in “Windows: Migrating a queue manager from a previous version to the latest version” on page 515. Apply the required maintenance level.
5. Start the cluster service on the selected node.
6. On the next node to be migrated, ensure that the IBM MQ resources are offline.
7. Remove this node from the list of possible owners. For clusters with more than two nodes, see the Additional considerations later in this topic.
8. Move the group containing the IBM MQ resource to one of the possible owners and bring it online.
9. Repeat steps 3-8 as necessary for any remaining nodes.

### Migrating a four-node MSCS cluster from an earlier version of the product to the latest version

The example in Table 82 on page 574 illustrates the steps involved in migrating a four-node MSCS cluster.

In the example IBM MQ resources include queue managers, applications, and dependant MSCS resources, such as an IP address defined as an MSCS resource. In each step, the changes are italicized.

**Step 1** Select the node to migrate and prepare it for upgrading from an earlier version of the product to the latest version.

1. Select node 1 to be migrated and convert it into a Passive node with no running IBM MQ resources.

2. Modify the possible owners of the group containing the IBM MQ resources, to encompass only the required online nodes. Failover does not attempt to switch IBM MQ resources to the node that is not a possible owner. It is safe to migrate that node.
3. Move the group containing the IBM MQ resource to one of the nodes that is a possible owner, and bring it online.
4. Stop the cluster service on the node being migrated. Stopping the service clears the MSCS cache of any IBM MQ libraries that have been registered for MSCS. The node goes offline.

**Step 2** Migrate IBM MQ from an earlier version of the product to the latest version

**Step 3** Start the cluster service on the selected node. The node becomes online, but it is not a possible owner, so no work is switched to it.

**Step 4** Repeat steps 1 - 3 for node 2. Nodes 1 and 2 are now online, and you have migrated them to the latest version. They are still doing no work, as they are not possible owners of any of the IBM MQ resource groups.

**Step 5** Migrate the cluster from running an earlier version of the product to the latest version. The number of migrated nodes is now greater or equal to the number of unmigrated nodes.

1. Change the set of possible owners from 3,4 to 1,2.
2. Move the IBM MQ resource groups from nodes 3 and 4 to nodes 1 and 2 and bring online.
3. From this point onward, the list of possible owners must include migrated nodes only. The IBM MQ resource must never failover to a node running a back level version of the product.

**Note:** If you must revert IBM MQ to an earlier version, the IBM MQ resources must be removed from MSCS control, before performing an uninstallation of IBM MQ

**Step 6** Migrate node 3 to the latest version.

1. Follow steps 1 - 3 for node 3.
2. Add node 3 to the list of possible owners.
3. Move the QMC resource group back from node 1 to node 3 and bring online again.

**Step 7** Repeat step 6 for node 4.

Table 82. Migrating a four-node MSCS cluster

| Steps  |         | 0               | 1               | 2                     | 3               | 4                     | 5               | 6                     | 7              |
|--------|---------|-----------------|-----------------|-----------------------|-----------------|-----------------------|-----------------|-----------------------|----------------|
| Node 1 | State   | Online          | <i>Offline</i>  | <i>Offline</i>        | <i>Online</i>   | Online                | Online          | Online                | Online         |
|        | Version | Earlier version | Earlier version | <i>Latest version</i> | Latest version  | Latest version        | Latest version  | Latest version        | Latest version |
|        | Groups  | QMA             |                 |                       |                 |                       | QMC,<br>QMA     | QMA                   | QMA            |
| Node 2 | State   | Online          | Online          | Online                | Online          | Online                | Online          | Online                | Online         |
|        | Version | Earlier version | Earlier version | Earlier version       | Earlier version | <i>Latest version</i> | Latest version  | Latest version        | Latest version |
|        | Groups  | QMB             | QMB             | QMB                   | QMB             |                       | QMD,<br>QMB     | QMD,<br>QMB           | QMB            |
| Node 3 | State   | Online          | Online          | Online                | Online          | Online                | Online          | Online                | Online         |
|        | Version | Earlier version | Earlier version | Earlier version       | Earlier version | Earlier version       | Earlier version | <i>Latest version</i> | Latest version |
|        | Groups  | QMC             | QMC,<br>QMA     | QMC,<br>QMA           | QMC,<br>QMA     | QMC,<br>QMA           |                 | QMC                   | QMC            |

Table 82. Migrating a four-node MSCS cluster (continued)

| Steps           |         | 0               | 1               | 2               | 3               | 4               | 5               | 6               | 7                     |
|-----------------|---------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------------|
| Node 4          | State   | Online          | Online          | Online          | Online          | Online          | Online          | Online          | Online                |
|                 | Version | Earlier version | Earlier version | Earlier version | Earlier version | Earlier version | Earlier version | Earlier version | <i>Latest version</i> |
|                 | Groups  | QMD             | QMD             | QMD             | QMD             | QMD, QMB        |                 |                 | QMD                   |
| Possible Owners |         | 1,2,3,4         | 2,3,4           | 2,3,4           | 2,3,4           | 3,4             | 1,2             | 1,2,3           | 1,2,3,4               |
| Task            |         |                 | Update 1        |                 |                 | Update 2        | Transfer        | Update 3        | Update 4              |

## What to do next

**Additional considerations in an MSCS setup with more than 2 nodes:** A cluster might contain enough nodes for you to form a group of migrated queue managers and a group of unmigrated nodes. Switch to the migrated group when it contains half the number of queue managers. Before you have reached the half way point, the unmigrated group are possible owners. When you reach the half way point, switch the possible owners to the migrated group.

### Related concepts:

“Migrating a queue manager in a high-availability configuration” on page 481

Follow standard procedures to migrate a queue manager that is part of a high-availability configuration. (On platforms other than z/OS.)

### Related reference:

“Windows: MSCS restriction with multiple installations” on page 621

When you install or upgrade to Version 7.1 or later, the first installation of the product on the server is the only one that can be used with Microsoft Cluster Server (MSCS). No other installations on the server can be used with MSCS. This restriction limits the use of MSCS with multiple installations of the product.

---

## Migrating from a single instance to a multi-instance queue manager

To migrate a single instance queue manager to a multi-instance queue manager, you must move the queue manager data to a shared directory, and reconfigure the queue manager on two other servers.

### Before you begin

You must check the prerequisites for running a multi-instance queue manager as part of this task. Some environments have been tested with multi-instance queue managers, and are known to work. They are AIX, Red Hat Linux, SUSE Linux Enterprise Server, HP-UX with the file system on Linux Red Hat, IBM i, and Windows Server. See Testing and support statement for IBM MQ multi-instance queue managers for the latest list of tested environments. The support statement has detailed version and prerequisite information for each environment it lists. Other environments might work; a test tool is provided with IBM MQ to assist you in qualifying other environments.

You must have three servers to run a multi-instance queue manager. One server has a shared file system to store the queue manager data and logs. The other servers run the active and standby instances of the queue manager.

### About this task

You have a single-instance queue manager that you want to convert to a multi-instance queue manager. The queue manager conversion itself is straightforward, but you must do other tasks to create a fully automated production environment.

You must check the prerequisites for a multi-instance queue manager, set up the environment and check it. You must set up a monitoring and management system to detect if the multi-instance queue manager has failed and been automatically restarted. You can then find out what caused the restart, remedy it, and restart the standby. You must also modify applications, or the way applications are connected to the queue manager, so that they can resume processing after a queue manager restart.

## Procedure

1. Check the operating system that you are going to run the queue manager on, and the file system on which the queue manager data and logs are stored on. Check that they can run a multi-instance queue manager.

- a. Consult Testing and support statement for IBM MQ multi-instance queue managers. See whether the combination of operating system and file system is capable of running a multi-instance queue manager.

A shared file system must provide lease-based locking to be adequate to run multi-instance queue managers. Lease-based locking is a recent feature of some shared file systems, and in some case fixes are required. The support statement provides you with the essential information.

- b. Run **amqmfscck** to verify that the file system is configured correctly.

File systems are sometimes configured with performance at a premium over data integrity. It is important to check the file system configuration. A negative report from the **amqmfscck** tool tells you the settings are not adequate. A positive result is an indication that the file system is adequate, but the result is not a definitive statement that the file system is adequate. It is a good indication.


- c. Run the integrity checking application provided in the technote, Testing a shared file system for compatibility with IBM MQ Multi-instance Queue Managers.

The checking application tests that the queue manager is restarting correctly.

2. Configure a user and group to be able to access a share on the networked file system from each server that is running a queue manager instance.

On Windows, the security IDs (SIDs) of the `mqm` group can be different; see Windows domains and multi-instance queue managers.

On UNIX and Linux, the `uid` and `gid` for `mqm` in `/etc/passwd` must be the same on each system; see Create a multi-instance queue manager on Linux .

 On IBM i, QMQM, QMQMADM, and any other user profiles that are granted access to the share must have the same passwords on all the servers.

3. Set up a directory for the share on the networked file system with the correct access permissions.


A typical configuration is to set up a single shared directory that contains all data and log directories for all queue managers that use the shared disk; see Share named `qmgrs` and log directories (Version 7.0.1 onwards) in Example directory configurations on UNIX systems.

For example, create a root directory on the share called `MQHA` that has subdirectories `data` and `logs`. Each queue manager creates its own data and log directories under `data` and `logs`. Create `MQHA` with the following properties:

On Windows, create `drive\MQHA` on the shared drive. The owner is a member of `mqm`. `mqm` must have full-control authority. Create a share for `drive\MQHA`.

On UNIX, create `/MQHA` on the shared drive. `/MQHA` is owned by the user and group `mqm` and has the access permissions `rwX`.


If you are using an NFS v4 file server, add the line `/MQHA * rw, sync, no_wdelay, fsid=0` to `etc/exports`, and then start the NFS daemon: `/etc/init.d/nfs start`.

 On IBM i, follow the instructions to create a network share using NetServer.

4. Copy the queue manager data and the logs to the share.



You can choose to copy files manually, by following the procedure to back up the queue manager. On Windows, you can run the **hamvmqm** command to move the queue manager data to the share. The **hamvmqm** command works for queue managers created before Version 7.0.1, and not reconfigured with a datapath, or for queue managers that do not have a **DataPath** configuration attribute. Choose one of these methods:

- Follow the instructions in Backing up queue manager data,  or Backups of IBM MQ for IBM i data, copying the queue manager data to the share. You must use this method if the **DataPath** configuration attribute is specified for this queue manager.
- Stop the queue manager, and then type the command,
 

```
hamvmqm /m /dd share\data /dd share\logs
```

Where *share* is to be the location of the data and logs that you created in step 3 on page 576.

5. Update the queue manager configuration information stored on the current queue manager server.

If you moved the queue manager data and logs by running the **hamvmqm** command, the command has already modified the configuration information correctly for you.

If you moved the queue manager data and logs manually, you must complete the following steps.

- On Windows:

- a. Modify the log registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\Configuration\QueueManager\QMGrName\LogPath="share\logs\QMGrName\"
```

- b. Modify the Prefix registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\Configuration\QueueManager\QMGrName\Prefix="share\data"
```

- On  IBM i, UNIX, and Linux,


- a. Modify the Log: stanza in the queue manager *qm.ini* file, which is on the *share*:

```
LogPath=share/logs/QMGrName
```

- b. Modify the QueueManager: stanza in the IBM MQ *mqs.ini* file, which is typically in the */var/mqm* directory on UNIX and Linux , and */QIBM/UserData/mqm* on IBM i:

```
DataPath=share/data/QMGrName
```

Where, *QMGrName* is the representation of the queue manager name in the existing registry key on Windows. *QMGrName* is the Directory name in the QueueManager: stanza in the *mqs.ini* file on

 IBM i, UNIX, and Linux. *share* is share where the data and logs are moved to.

6. Add the queue manager configuration information to the new queue manager server.

- a. Run the **dspmqinf** command to display the queue manager information

Run the command on the server that ran the queue manager in Version 6.0.

```
dspmqinf -o command QMGrName
```

The command output is formatted ready to create a queue manager configuration. **addmqinf -s QueueManager -v Name= QMGrName -v Directory= QMGrName -v Prefix=d:\var\mqm Datapath= \share\data\QMGrName**

- b. Create a queue manager configuration on the other server.

Run the **addmqinf** command copied from the previous output

7. Add the network address of the new server to the connection name in client and channel definitions.

- a. Find all the client, sender, and requester TCPIP settings that refer to the server.

Client settings might be in Client Definition Tables (CCDT), in environment variables, in Java properties files, or in client code.

Cluster channels automatically discover the connection name of a queue manager from its cluster receiver channel. As long as the cluster receiver channel name is blank or omitted, TCPIP discovers the IP address of the server hosting the queue manager.

- b. Modify the connection name for each of these connections to include the TCPIP addresses of both servers that are hosting the multi-instance queue manager.

For example, change:

```
echo DISPLAY CHANNEL(ENGLAND) CONNAME | runmqsc QM1
5724-H72 (C) Copyright IBM Corp. 1994, 2009. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM1.
1: DISPLAY CHANNEL(ENGLAND) CONNAME
AMQ8414: Display Channel details.
CHANNEL(ENGLAND) CHLTYPE(SDR)
CONNAME(LONDON)
```

Into:

```
echo ALTER CHANNEL(ENGLAND) CHLTYPE(SDR) CONNAME('LONDON, BRISTOL') | runmqsc QM1
```

8. Update your monitoring and management procedures to detect the queue manager restarting.
9. Update client applications to be automatically reconnectable, if appropriate.
10. Update the start procedure for your IBM MQ applications to be started as queue manager services.
11. Start each instance of the queue manager, permitting them to be highly available.

The first instance of the queue manager that is started becomes the active instance.

Issue the command twice, once on each server.

```
strmqm -x QMgrName
```

## What to do next

To get the highest availability out of multi-instance queue managers, you must design client applications to be reconnectable and server applications to be restartable; see Application recovery.

### Related information:

**amqmfsc** (file system check)

Application recovery

Automatic client reconnection

Backing up queue manager data


Channel and client reconnection

Changing configuration information on Windows, UNIX and Linux systems

Create a multi-instance queue manager on Linux


Moving a queue manager to MSCS storage

Multi-instance queue managers

 Multi-instance queue managers on IBM i

Queue manager configuration files, qm.ini

Shared file system

 [Testing a shared file system for compatibility with IBM MQ Multi-instance Queue Managers](#)

 [Testing and support statement for IBM MQ multi-instance queue managers](#)

The IBM MQ configuration file, mqs.ini

 The IBM MQ configuration file mqs.ini - IBM i

Verifying shared file system locking

Windows domains and multi-instance queue managers

Working with services

---

## Reverting to a single-instance queue manager

Revert a multi-instance queue manager to a single instance queue manager by stopping the standby instance. Then restart the active instance and do not set the flag that permits standby instances.

### Before you begin

You have at least three servers configured to run a queue manager as a multi-instance queue manager. The queue manager is currently running as a multi-instance queue manager, with one standby instance active.

### About this task

The task involves deactivating the active standby so that only the running multi-instance queue manager remains active. To prevent a standby instance being started in the future, you must stop the active instance and restart it. When you restart it, you start it as a single instance queue manager that prevents standby instances being started. The standby instance is stopped as a separate step, to give you the option of restarting the active instance at a later date. You can stop both instances by running the standard `endmqm QMgrName` command on the server running the active queue manager.

### Procedure

1. Stop the standby queue manager instance.

On the server running the standby instance:

- Windows, UNIX, and Linux

```
endmqm -w QMgrName
```

-  IBM i

```
ENDMQM MQMNAME (QMgrName) *WAIT
```

2. Stop the active queue manager instance.

On the server running the active instance:

- Windows, UNIX, and Linux

```
endmqm -w (QMgrName)
```

-  IBM i

```
ENDMQM MQMNAME (QMgrName) *WAIT
```

3. Restart the queue manager, preventing standbys.

On the server that is going to run the queue manager:

- Windows, UNIX, and Linux

```
strmqm QMgrName
```

-  IBM i

```
STRMQM MQMNAME (QMgrName)
```

### What to do next

You might want to run the queue manager as a single instance on the same server as the queue manager data.

When the queue manager is stopped move the queue manager data back to the server that is running the queue manager. Alternatively install IBM MQ, and then move the queue manager configuration definition onto the server with the queue manager data. Both tasks are variations of steps in “Migrating from a single instance to a multi-instance queue manager” on page 575 to create a multi-instance queue manager.

---

## Changes that affect migration

Changes are listed that affect the migration of a queue manager to the current release of IBM MQ, or that affect existing applications or configurations.

Review the list of changes before upgrading queue managers to the latest product version. Decide whether you must plan to make changes to existing applications, scripts, and procedures before starting to migrate your systems.

### Note:

**distributed** **IBM i** On distributed platforms, you cannot reverse queue manager migration to remove the effect of changes.

**z/OS** On z/OS, you can reverse queue manager migration as long as you have not enabled new function. You enable new function by setting the **OPMODE** parameter to (NEWFUNC, 800).

## Changes for IBM MQ Version 8.0

- Command changes: **dmpmqcfg** output
- Command changes: results displayed in response to DISPLAY CONN command from **runmqsc**
- Command level: changes
- **Solaris** Deprecation: linking with libmqmcs and libmqmzse libraries
- Java and JMS: Changes to CipherSuite support
- Java: changes for Java 7
- Java: Changes to IBM MQ classes for Java
- Java: Changes to MQException and MQDataException class logging
- **V8.0.0.3** Extended start events for multi-instance queue managers
- JMS: changes to JAR files for JMS 2.0
- JMS: changes to IBM MQ classes for JMS
- Publish/subscribe: changes to DISPLAY TOPIC and DISPLAY TPSTATUS output
- Publish/subscribe: changes to how the subscription selection string is evaluated
- Queue managers: change to default TCP buffer size
- Queue managers: Version attribute added for queue managers in a cluster
- Security: new CONNAUTH CHCKLOCL parameter
- Security: Disable AMS at the client
- Security: reduction in channel send exit buffer space for SSL and TLS
- Solaris client: Change in method of installation
- **z/OS** z/OS: changes to log RBA and URID lengths
- **z/OS** z/OS: WLM/DNS no longer supported


For details of all changes in this release, including those that affect migration, see What's changed in IBM MQ Version 8.0 .

### Related concepts:

“Introduction to IBM MQ migration” on page 421

These topics explain the concepts that you must understand before planning migration tasks, where to find migration topics, and which migration paths are supported.

## Coexistence

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations. In addition to queue managers coexisting on a server, objects, and commands must work correctly with different queue managers running at different command levels.

 z/OS

### Multiple queue manager versions in z/OS


There can be several IBM MQ subsystems in a z/OS image, and they can use different versions of IBM MQ, provided that the IBM MQ early code modules are of the latest version being used. (These modules are loaded at z/OS IPL time and are shared among all the IBM MQ subsystems in the z/OS image.)

This means that you can run one queue manager at the latest version and another in the same image with an earlier version, provided that the early code is that of the latest version.

The coexistence section lists restrictions in the use of objects and commands when they are used with queue managers at multiple command levels. The queue managers might be running on a single server, or in a cluster.

### Related concepts:

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

### Related tasks:

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

## JMS: Administered objects

Version 6.0 JMS administered objects, such as connection factories and destinations are compatible with later releases.

JMS applications can use connection factory or destination objects created in IBM WebSphere MQ Version 6.0. Any new properties, that did not exist in Version 6.0, assume their default value.

IBM WebSphere MQ Version 6.0 JMS applications can use connection factory or destination objects created in later releases. Any of the new properties that do not exist in Version 6.0 are ignored.

### **Mixed version cluster coexistence**

A cluster can contain queue managers running at IBM MQ Version 8.0, and any currently supported level of IBM MQ. However new features cannot be exploited from queue managers at an earlier level.

### **Cluster workload balancing in a mixed version cluster**

IBM MQ Version 7.1 added a new **DEFBIND** value called **GROUP** to queues. Applications on queue managers earlier than Version 7.1 must not open or put messages to queues specifying the new value. When an application ignores this limitation, the workload balancing behavior (for example: **BIND\_ON\_OPEN** or **BIND\_NOT\_FIXED**) is undefined.

### **Routing behavior in a mixed version publish/subscribe cluster**

As described in What's new in IBM MQ Version 8.0 , this version adds topic host routing for publish/subscribe clusters. As for all new clustered behavior, the queue manager where the object is defined, and the full repository queue managers, must be at a level that supports the new feature. Any queue manager in the cluster that is at an earlier level does not adhere to the new behavior.

When a clustered topic is defined for topic host routing (by setting the topic **CLROUTE** parameter to **TOPICHOST** ), only queue managers at the new level are aware of the clustered topic. Older queue managers do not receive the clustered topic definition and therefore behave as if the topic is not clustered. This means that all queue managers that need to work in a routed publish/subscribe manner must be at a version that supports this feature, not just the queue managers that host the routed topics.

#### **Important notes:**

- All full repositories must be at Version 8.0 to use this feature. If a full repository queue manager is at an earlier version, the **CLROUTE** of **TOPICHOST** is not recognized by the full repository, and the full repository propagates the topic definition to all queue managers in the cluster. Any pre- Version 8.0 queue managers then use the topic as if it is defined for **DIRECT** routing. This behavior is unsupported.
- If an older queue manager defines a direct routed clustered topic with the same name as an existing topic host routed clustered topic, the full repositories notice the conflicting definition and do not propagate the definition.

To find out the version of each queue manager in the cluster, specify the **VERSION** parameter with the **DISPLAY CLUSQMGR** command. If you issue this command from a queue manager with a full repository, the information returned applies to every queue manager in the cluster. Otherwise the information returned applies only to the queue managers in which it has an interest. That is, every queue manager to which it has tried to send a message and every queue manager that holds a full repository.

## z/OS: ISPF operations and control panels

When using the operations and control panels, the IBM MQ libraries you use in ISPF must be compatible with the queue manager you are working with.

Table 83 shows which versions of the operations and controls panels you use in ISPF are compatible with which levels of queue manager. Version 7.1 panels are incompatible with any release before Version 6.0.

Table 83. Compatibility of queue manager versions with operations and control panel versions

| Version             | Queue-sharing group containing a mixture of Version 7.0.1, Version 7.1, and Version 8.0 queue managers. | Version 8.0 queue manager                 | Version 7.1 queue manager                 | Version 7.0.1 queue manager               | Version 7.0 queue manager |
|---------------------|---------------------------------------------------------------------------------------------------------|-------------------------------------------|-------------------------------------------|-------------------------------------------|---------------------------|
| Version 8.0 panel   | Compatible with restrictions and warnings                                                               | Compatible                                | Compatible with restrictions and warnings |                                           |                           |
| Version 7.1 panel   | Not compatible                                                                                          | Compatible                                |                                           | Compatible with restrictions and warnings |                           |
| Version 7.0.1 panel | Not compatible                                                                                          | Compatible with restrictions and warnings |                                           | Compatible                                |                           |
| Version 7.0 panel   | Not compatible                                                                                          | Not compatible                            | Compatible with restrictions and warnings | Compatible                                |                           |

## z/OS: Queue sharing group coexistence



A queue-sharing group can contain queue managers running on IBM WebSphere MQ Version 7.0.1, and on later releases. The queue managers can access the same shared queues and other shared objects. Queue managers running earlier versions of the product must have the coexistence PTF applied for the latest release.

**Note:** After the coexistence PTF has been applied, the earlier version queue managers must be started at least once.

A queue-sharing group containing a latest version queue manager, can contain queue managers only at one other, earlier version. You can not start a latest version queue manager in a queue-sharing group that contains other queue managers at more than one other version.

Only run queue managers in a mixed-version queue-sharing group for the time it takes to migrate all queue managers to the later version. If the queue-sharing group contains queue managers with a mixture of versions, new functions on the latest version, that are restricted by OPMODE, will not be available.

## z/OS: Properties of objects in a mixed queue-sharing group

Attributes that did not exist in earlier versions can be created and altered on queue managers of a later version in a mixed queue-sharing group. The attributes are not available to queue managers in the group that are at an earlier level.

Any **QSGDISP**(GROUP) TOPIC objects having the **CLROUTE**(TOPICHOST) attribute set, and any **QSGDISP**(GROUP) AUTHINFO objects with **AUTHTYPE**(IDPWOS), are hidden from queue managers earlier than Version 8.0 in a mixed queue-sharing group.

## z/OS: MQSC commands in a mixed queue-sharing group



Existing **MQSC** commands using new keywords and attribute values can be entered for routing to a migrated queue manager. You can enter the commands on any queue manager. Route the commands using **CMDSCOPE**. Commands with new keywords and attribute values, or new commands, routed to a previous version of queue manager, fail.


## Changes from WebSphere MQ Version 7.0.1 to IBM MQ Version 8.0

All IBM MQ changes that affect the migration of a queue manager from Version 7.0.1 to Version 8.0 are listed.

New capabilities that do not affect existing IBM MQ applications are not listed.

### Note:

  On distributed platforms, you cannot reverse queue manager migration to remove the effect of changes.

 On z/OS, you can reverse queue manager migration as long as you have not enabled new function. You enable new function by setting the **OPMODE** parameter to (NEWFUNC, 800).

Click on a link to read the details about the change.

### Related reference:

 [“z/OS: OPMODE” on page 677](#)

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. To access V8.0 capabilities, change the value of **OPMODE** to **OPMODE=(NEWFUNC,800)**. To restrict the use of new capabilities, and retain the ability to revert the queue manager to its earlier level, leave **OPMODE** at its default setting, **OPMODE=(COMPAT,800)**.

## Version 7.0.1 features and their supported APIs

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

The tables in this section show supported IBM WebSphere MQ Version 7.0.1 APIs. Use the information in these tables in conjunction with the information in the system requirements web pages. For links to system requirements information for all releases of IBM WebSphere MQ or IBM MQ, see the IBM MQ System Requirements web page.

From IBM MQ Version 8.0, you can use the Software Product Compatibility Reports (SPCR) tool to find information on supported operating systems, system requirements, prerequisites, and optional supported software. For more information about the SPCR tool and links to reports for each supported platform, see the System Requirements for IBM MQ Version 8.0 web page.



## Version 7.0.1 features and their supported APIs: JMS:

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

### JMS

The following tables show which features and environments are supported for JMS.

Table 84. JMS: API support

|                                     |                    | JMS                                |                |                     |                |                  |               |                     |                |      |     |  |
|-------------------------------------|--------------------|------------------------------------|----------------|---------------------|----------------|------------------|---------------|---------------------|----------------|------|-----|--|
|                                     |                    | Distributed platforms              |                |                     |                | z/OS             |               |                     |                |      |     |  |
| API                                 | Basic <sup>1</sup> | WAS and other Java EE <sup>2</sup> | Other TM       | Db2 SP <sup>3</sup> | Basic          | WAS <sup>2</sup> | Other Java EE | Db2 SP <sup>3</sup> | Other TM       | CICS | IMS |  |
| V7 Publish/Subscribe                | Y                  | Y                                  | Y              | N/A <sup>4</sup>    | Y              | Y                | N/A           | N/A <sup>4</sup>    |                | N/A  | N/A |  |
| Browse with Mark                    | N                  | N                                  | N              | N/A <sup>4</sup>    | N              | N                | N/A           | N/A <sup>4</sup>    |                | N/A  | N/A |  |
| Message Properties                  | Y <sup>5</sup>     | Y <sup>5</sup>                     | Y <sup>5</sup> | N/A <sup>4</sup>    | Y <sup>5</sup> | Y <sup>5</sup>   | N/A           | N/A <sup>4</sup>    | Y <sup>5</sup> | N/A  | N/A |  |
| Message Selectors                   | Y                  | Y                                  | Y              | N/A <sup>4</sup>    |                |                  | N/A           | N/A <sup>4</sup>    |                | N/A  | N/A |  |
| Asynchronous consume                | Y                  | Y                                  | Y              | N/A <sup>4</sup>    |                |                  | N/A           | N/A <sup>4</sup>    |                | N/A  | N/A |  |
| Content filtering Publish/Subscribe | N                  | N                                  | N              | N/A <sup>4</sup>    | N              | N                | N/A           | N/A <sup>4</sup>    | N              | N/A  | N/A |  |
| Group/Segment messages              | N                  | N                                  | N              | N/A <sup>4</sup>    | N              | N                | N/A           | N/A <sup>4</sup>    | N              | N/A  | N/A |  |
| PCF Classes                         | Y <sup>6</sup>     | Y <sup>6</sup>                     | Y <sup>6</sup> | N/A <sup>4</sup>    | Y <sup>6</sup> | Y <sup>6</sup>   | N/A           | N/A <sup>4</sup>    | Y <sup>6</sup> | N/A  | N/A |  |
| Global TX Participant               | Y <sup>7</sup>     | Y                                  | Y <sup>7</sup> | N/A <sup>4</sup>    | Y <sup>8</sup> | Y                | N/A           | N/A <sup>4</sup>    | Y <sup>8</sup> | N/A  | N/A |  |
| Global TX Coordinator               | N                  | N                                  | N              | N/A <sup>4</sup>    | N              | N                | N/A           | N/A                 | N              | N/A  | N/A |  |

#### Note:

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Db2 SP Distributed: cannot work reliably because of the restricted Java environment.
5. Cannot see the whole MQI namespace, therefore an MQI application might generate messages whose properties cannot be read.

6. Can build/parse messages but not use the MessageAgent classes and methods.
7. Open source JTA coordinators might not be supported.
8. RRS is the coordinator.

Table 85. JMS: Client communication (not running on top of C client)

| Client communication (not running on top of C client) | JMS                   |                                    |                |                     |       |                  |               |                     |          |      |     |
|-------------------------------------------------------|-----------------------|------------------------------------|----------------|---------------------|-------|------------------|---------------|---------------------|----------|------|-----|
|                                                       | Distributed platforms |                                    |                |                     | z/OS  |                  |               |                     |          |      |     |
|                                                       | Basic <sup>1</sup>    | WAS and other Java EE <sup>2</sup> | Other TM       | Db2 SP <sup>3</sup> | Basic | WAS <sup>2</sup> | Other Java EE | Db2 SP <sup>3</sup> | Other TM | CICS | IMS |
| Read ahead                                            | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| Asynchronous output                                   | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| Pre-connect exit                                      | N                     | N                                  | N              | N/A                 | N/A   | N                | N/A           | N/A                 | N/A      | N/A  | N/A |
| CCDT                                                  | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| Multiple CONNAME                                      | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| Auto reconnect                                        | Y                     | Y <sup>5</sup>                     | Y              | N/A                 | N/A   | Y <sup>5</sup>   | N/A           | N/A                 | N/A      | N/A  | N/A |
| Channel compression                                   | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| Exit chaining                                         | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |
| MQCSP/<br>MQXR_SEC_PARMS                              | Y <sup>4</sup>        | Y <sup>4</sup>                     | Y <sup>4</sup> | N/A                 | N/A   | Y <sup>4</sup>   | N/A           | N/A                 | N/A      | N/A  | N/A |
| Native SSL                                            | Y                     | Y                                  | Y              | N/A                 | N/A   | Y                | N/A           | N/A                 | N/A      | N/A  | N/A |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Not consistent across all clients.
5. Use RA features for reconnect, not the FAP-level operation.

**Related information:**

[IBM MQ resource adapter V8.0 statement of support](#)

Running IBM MQ classes for Java applications within Java Platform, Enterprise Edition

**Version 7.0.1 features and their supported APIs: Java:**

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

**Java**

The following tables show which features and environments are supported for Java.

*Table 86. Java: API support*

| API                                 | Base Java             |                                    |          |                     |                  |                                    |                     |          |      |     |
|-------------------------------------|-----------------------|------------------------------------|----------|---------------------|------------------|------------------------------------|---------------------|----------|------|-----|
|                                     | Distributed platforms |                                    |          |                     | z/OS             |                                    |                     |          |      |     |
|                                     | Basic <sup>1</sup>    | WAS and other Java EE <sup>2</sup> | Other TM | Db2 SP <sup>3</sup> | Basic            | WAS and other Java EE <sup>2</sup> | Db2 SP <sup>3</sup> | Other TM | CICS | IMS |
| V7 Publish/Subscribe                | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | Y                | N/A                                | N/A <sup>6</sup>    | N/A      | Y    | N/A |
| Browse with Mark                    | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | Y                | N/A                                | N/A <sup>6</sup>    | N/A      | Y    | N/A |
| Message Properties                  | N                     | N/A                                | N/A      | N/A <sup>6</sup>    | N                | N/A                                | N/A <sup>6</sup>    | N/A      | N    | N/A |
| Message Selectors                   | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | Y                | N/A                                | N/A <sup>6</sup>    | N/A      | Y    | N/A |
| Asynchronous consume                | N                     | N/A                                | N/A      | N/A <sup>6</sup>    | N                | N/A                                | N/A <sup>6</sup>    | N/A      | N    | N/A |
| Content filtering Publish/Subscribe | N                     | N/A                                | N/A      | N/A <sup>6</sup>    | N                | N/A                                | N/A <sup>6</sup>    | N/A      | N    | N/A |
| Group/Segment messages              | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | Y                | N/A                                | N/A <sup>6</sup>    | N/A      | Y    | N/A |
| PCF Classes                         | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | Y                | N/A                                | N/A <sup>6</sup>    | N/A      | Y    | N/A |
| Global TX Participant               | N                     | N/A                                | N/A      | N/A <sup>6</sup>    | N                | N/A                                | N/A <sup>6</sup>    | N/A      | N    | N/A |
| Global TX Coordinator               | Y                     | N/A                                | N/A      | N/A <sup>6</sup>    | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.

3. Db2 Stored Procedures.
4. While it might work (some customers have used this environment), it is not recommended or fully supported.
5. RRS is the coordinator.
6. Db2 SP on Distributed platform cannot work reliably because of the restricted Java environment.

Table 87. Java: Client communication (not running on top of C client)

| Client communication (not running on top of C client) | Base Java             |                                    |          |                     |                  |                                    |                     |          |      |     |
|-------------------------------------------------------|-----------------------|------------------------------------|----------|---------------------|------------------|------------------------------------|---------------------|----------|------|-----|
|                                                       | Distributed platforms |                                    |          |                     | z/OS             |                                    |                     |          |      |     |
|                                                       | Basic <sup>1</sup>    | WAS and other Java EE <sup>2</sup> | Other TM | Db2 SP <sup>3</sup> | Basic            | WAS and other Java EE <sup>2</sup> | Db2 SP <sup>3</sup> | Other TM | CICS | IMS |
| Read ahead                                            | Y                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Asynchronous output                                   | N                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Pre-connect exit                                      | N                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| CCDT                                                  | Y                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Multiple CONNAME                                      | N                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Auto reconnect                                        | N                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Channel compression                                   | Y                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Exit chaining                                         | Y                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| MQCSP/MQXR_SEC_PARMS                                  | Y <sup>4</sup>        | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |
| Native SSL                                            | Y                     | N/A                                | N/A      | N/A                 | N/A <sup>5</sup> | N/A                                | N/A                 | N/A      | N/A  | N/A |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Not consistent across all clients.
5. Java client not supported in any environment on z/OS

**Related information:**

 [IBM MQ resource adapter V8.0 statement of support](#)

Running IBM MQ classes for Java applications within Java Platform, Enterprise Edition

**Version 7.0.1 features and their supported APIs: XMS .NET:**

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

**XMS .NET**

The following tables show which features and environments are supported for XMS .NET.

*Table 88. XMS .NET: API support*

| API                                    | XMS .NET              |                  |               |          |                     |
|----------------------------------------|-----------------------|------------------|---------------|----------|---------------------|
|                                        | Distributed platforms |                  |               |          |                     |
|                                        | Basic <sup>1</sup>    | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| V7 Publish<br>Subscribe                | Y                     | N/A              | N/A           | N/A      | N/A                 |
| Browse with<br>Mark                    | N                     | N/A              | N/A           | N/A      | N/A                 |
| Message<br>Properties                  | Y <sup>4</sup>        | N/A              | N/A           | N/A      | N/A                 |
| Message Selectors                      | Y                     | N/A              | N/A           | N/A      | N/A                 |
| Asynchronous<br>consume                | Y                     | N/A              | N/A           | N/A      | N/A                 |
| Content filtering<br>Publish/Subscribe | N                     | N/A              | N/A           | N/A      | N/A                 |
| Group/Segment<br>messages              | N                     | N/A              | N/A           | N/A      | N/A                 |
| PCF Classes                            | N                     | N/A              | N/A           | N/A      | N/A                 |
| Global TX<br>Participant               | Y <sup>5</sup>        | N/A              | N/A           | N/A      | N/A                 |
| Global TX<br>Coordinator               | N                     | N/A              | N/A           | N/A      | N/A                 |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Cannot see the whole MQI namespace, therefore an MQI application might generate messages whose properties cannot be read.
5. Works with Microsoft DTC only.

Table 89. XMS .NET: Client communication (not running on top of C client)

|                                                       | XMS .NET           |                  |               |          |                     |
|-------------------------------------------------------|--------------------|------------------|---------------|----------|---------------------|
|                                                       | Distributed        |                  |               |          |                     |
| Client communication (not running on top of C client) | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| Read ahead                                            | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Asynchronous put                                      | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Pre-connect exit                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| CCDT                                                  | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Multiple CONNAME                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| Auto reconnect                                        | N                  | N/A              | N/A           | N/A      | N/A                 |
| Channel compression                                   | N                  | N/A              | N/A           | N/A      | N/A                 |
| Exit chaining                                         | N                  | N/A              | N/A           | N/A      | N/A                 |
| MQCSP/<br>MQXR_SEC_PARMS                              | N                  | N/A              | N/A           | N/A      | N/A                 |
| Native SSL                                            | N                  | N/A              | N/A           | N/A      | N/A                 |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.

**Version 7.0.1 features and their supported APIs: XMS C and C++:**

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

**XMS C and C++**

The following tables show which features and environments are supported for XMS C and C++.

Table 90. XMS C and C++: API support

|                         | XMS C and C++      |                  |               |          |                     |
|-------------------------|--------------------|------------------|---------------|----------|---------------------|
|                         | Distributed        |                  |               |          |                     |
| API                     | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| V7 Publish<br>Subscribe | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Browse with<br>Mark     | N                  | N/A              | N/A           | N/A      | N/A                 |
| Message<br>Properties   | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Message Selectors       | Y                  | N/A              | N/A           | N/A      | N/A                 |

Table 90. XMS C and C++: API support (continued)

|                                     | XMS C and C++      |                  |               |          |                     |
|-------------------------------------|--------------------|------------------|---------------|----------|---------------------|
|                                     | Distributed        |                  |               |          |                     |
| API                                 | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| Asynchronous consume                | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Content filtering Publish/Subscribe | N                  | N/A              | N/A           | N/A      | N/A                 |
| Group/Segment messages              | N                  | N/A              | N/A           | N/A      | N/A                 |
| PCF Classes                         | N                  | N/A              | N/A           | N/A      | N/A                 |
| Global TX Participant               | N                  | N/A              | N/A           | N/A      | N/A                 |
| Global TX Coordinator               | N                  | N/A              | N/A           | N/A      | N/A                 |

**Note:**

1. Also OSGL.
2. WebSphere Application Server.
3. Db2 Stored Procedures.

Table 91. XMS C and C++: Client communication (not running on top of C client)

|                                                       | XMS C and C++      |                  |               |          |                     |
|-------------------------------------------------------|--------------------|------------------|---------------|----------|---------------------|
|                                                       | Distributed        |                  |               |          |                     |
| Client communication (not running on top of C client) | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| Read ahead                                            | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Asynchronous put                                      | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Pre-connect exit                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| CCDT                                                  | N                  | N/A              | N/A           | N/A      | N/A                 |
| Multiple CONNAME                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| Auto reconnect                                        | N                  | N/A              | N/A           | N/A      | N/A                 |
| Channel compression                                   | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Exit chaining                                         | Y                  | N/A              | N/A           | N/A      | N/A                 |
| MQCSP/MQXR_SEC_PARMS                                  | N/A                | N/A              | N/A           | N/A      | N/A                 |
| Native SSL                                            | Y                  | N/A              | N/A           | N/A      | N/A                 |

**Note:**

1. Also OSGL.
2. WebSphere Application Server.
3. Db2 Stored Procedures.

## Version 7.0.1 features and their supported APIs: .NET:

Use this information to learn about the IBM WebSphere MQ Version 7.0.1 APIs with features and environments that might not be as fully supported as the C MQI.

### .NET

The following table shows which features and environments are supported for .NET.

Table 92. .NET: API support

| API                                 | .NET               |                  |               |          |                     |
|-------------------------------------|--------------------|------------------|---------------|----------|---------------------|
|                                     | Distributed        |                  |               |          |                     |
|                                     | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| V7 Publish Subscribe                | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Browse with Mark                    | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Message Properties                  | Y <sup>4</sup>     | N/A              | N/A           | N/A      | N/A                 |
| Message Selectors                   | N                  | N/A              | N/A           | N/A      | N/A                 |
| Asynchronous consume                | N                  | N/A              | N/A           | N/A      | N/A                 |
| Content filtering Publish/Subscribe | N                  | N/A              | N/A           | N/A      | N/A                 |
| Group/Segment messages              | N                  | N/A              | N/A           | N/A      | N/A                 |
| PCF Classes                         | N                  | N/A              | N/A           | N/A      | N/A                 |
| Global TX Participant               | Y <sup>5</sup>     | N/A              | N/A           | N/A      | N/A                 |
| Global TX Coordinator               | N                  | N/A              | N/A           | N/A      | N/A                 |

#### Note:

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Cannot see the whole MQI namespace, therefore an MQI application might generate messages whose properties cannot be read.
5. Works with Microsoft DTC only.



Table 93. .NET: Client communication (not running on top of C client)

| Client communication (not running on top of C client) | .NET               |                  |               |          |                     |
|-------------------------------------------------------|--------------------|------------------|---------------|----------|---------------------|
|                                                       | Distributed        |                  |               |          |                     |
|                                                       | Basic <sup>1</sup> | WAS <sup>2</sup> | Other Java EE | Other TM | Db2 SP <sup>3</sup> |
| Read ahead                                            | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Asynchronous put                                      | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Pre-connect exit                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| CCDT                                                  | Y                  | N/A              | N/A           | N/A      | N/A                 |
| Multiple CONNAME                                      | N                  | N/A              | N/A           | N/A      | N/A                 |
| Auto reconnect                                        | N                  | N/A              | N/A           | N/A      | N/A                 |
| Channel compression                                   | N                  | N/A              | N/A           | N/A      | N/A                 |
| Exit chaining                                         | N                  | N/A              | N/A           | N/A      | N/A                 |
| MQCSP/<br>MQXR_SEC_PARMS                              | N                  | N/A              | N/A           | N/A      | N/A                 |
| Native SSL                                            | N <sup>4</sup>     | N/A              | N/A           | N/A      | N/A                 |

**Note:**

1. Also OSGI.
2. WebSphere Application Server.
3. Db2 Stored Procedures.
4. Support available in Unmanaged mode only.

**Channel authentication**

From IBM WebSphere MQ Version 7.1 onwards, when you migrate a queue manager from an earlier release, channel authentication using channel authentication records is disabled. Channels continue to work as before. If you create a queue manager in Version 8.0, channel authentication using channel authentication records is enabled, but with minimal additional checking. Some channels might fail to start.

**Migrated queue managers**

Channel authentication is disabled for migrated queue managers.

To start using channel authentication records you must run this MQSC command:

```
ALTER QMGR CHLAUTH(ENABLED)
```

**New queue managers**

Channel authentication is enabled for new queue managers.

You want to connect existing queue managers or IBM MQ MQI client applications to a newly created queue manager. Most connections work without specifying any channel authentication records. The following exceptions are to prevent privileged access to the queue manager, and access to system channels.

1. Privileged user IDs asserted by a client-connection channel are blocked by means of the special value \*MQADMIN.

```
SET CHLAUTH('*') TYPE(BLOCKUSER) USERLIST('*MQADMIN') +
DESCR('Default rule to disallow privileged users')
```

2. Except for the channel used by MQ Explorer, all SYSTEM.\* channels are blocked.

```
SET CHLAUTH('SYSTEM.*') TYPE(ADDRESSMAP) ADDRESS('*') USERSRC(NOACCESS) +
DESCR('Default rule to disable all SYSTEM channels')
```

```
SET CHLAUTH(SYSTEM.ADMIN.SVRCONN) TYPE(ADDRESSMAP) ADDRESS('*') USERSRC(CHANNEL) +
DESCR('Default rule to allow MQ Explorer access')
```

**Note:** This behavior is default for all new Version 8.0 queue managers on startup.

If you must work around the exceptions, you can run an MQSC command to add in more rules to allow channels blocked by the default rules to connect, or disable channel authentication checking:

```
ALTER QMGR CHLAUTH(DISABLED)
```

**Related information:**

Channel authentication records

### Change in behavior of the endmqm command

From IBM WebSphere MQ Version 7.5 onwards, issuing an **endmqm** command and **dspmq** command immediately after each other might return misleading status.

When issuing an **endmqm -c** or **endmqm -w** command, in the unlikely event that a **dspmq** command is issued in the small timeframe between the applications disconnecting and the queue manager actually stopping, the **dspmq** command might report the status as Ending immediately, even though a controlled shutdown is actually happening.

**Related information:**

endmqm

dspmq

### Changes to cluster error recovery (on servers other than z/OS )

From IBM WebSphere MQ Version 7.1 onwards, the queue manager reruns operations that caused problems, until the problems are resolved. If, after five days, the problems are not resolved, the queue manager shuts down to prevent the cache becoming more out of date.

Before Version 7.1, if a queue manager detected a problem with the local repository manager managing a cluster, it updated the error log. In some cases, it then stopped managing clusters. The queue manager continued to exchange applications messages with a cluster, relying on its increasingly out of date cache of cluster definitions. From Version 7.1 onwards, the queue manager reruns operations that caused problems, until the problems are resolved. If, after five days, the problems are not resolved, the queue manager shuts down to prevent the cache becoming more out of date. As the cache becomes more out of date, it causes a greater number of problems. The changed behavior regarding cluster errors in Version 7.1 or later does not apply to z/OS.

Every aspect of cluster management is handled for a queue manager by the local repository manager process, amqrrmf. The process runs on all queue managers, even if there are no cluster definitions.

Before Version 7.1, if the queue manager detected a problem in the local repository manager, it stopped the repository manager after a short interval. The queue manager kept running, processing application messages and requests to open queues, and publish or subscribe to topics.

With the repository manager stopped, the cache of cluster definitions available to the queue manager became more out of date. Over time, messages were routed to the wrong destination, and applications failed. Applications failed attempting to open cluster queues or publication topics that had not been propagated to the local queue manager.

Unless an administrator checked for repository messages in the error log, the administrator might not realize the cluster configuration had problems. If the failure was not recognized over an even longer time, and the queue manager did not renew its cluster membership, even more problems occurred. The instability affected all queue managers in the cluster, and the cluster appeared unstable.

From Version 7.1 onwards, IBM MQ takes a different approach to cluster error handling. Rather than stop the repository manager and keep going without it, the repository manager reruns failed operations. If the queue manager detects a problem with the repository manager, it follows one of two courses of action.

1. If the error does not compromise the operation of the queue manager, the queue manager writes a message to the error log. It reruns the failed operation every 10 minutes until the operation succeeds. By default, you have five days to deal with the error; failing which, the queue manager writes a message to the error log, and shuts down. You can postpone the five day shutdown.
2. If the error compromises the operation of the queue manager, the queue manager writes a message to the error log, and shuts down immediately.

An error that compromises the operation of the queue manager is an error that the queue manager has not been able to diagnose, or an error that might have unforeseeable consequences. This type of error often results in the queue manager writing an FFST file. Errors that compromise the operation of the queue manager might be caused by a bug in IBM MQ, or by an administrator, or a program, doing something unexpected, such as ending an IBM MQ process.

The point of the change in error recovery behavior is to limit the time the queue manager continues to run with a growing number of inconsistent cluster definitions. As the number of inconsistencies in cluster definitions grows, the chance of abnormal application behavior grows with it.

The default choice of shutting down the queue manager after five days is a compromise between limiting the number of inconsistencies and keeping the queue manager available until the problems are detected and resolved.

You can extend the time before the queue manager shuts down indefinitely, while you fix the problem or wait for a planned queue manager shutdown. The five-day stay keeps the queue manager running through a long weekend, giving you time to react to any problems or prolong the time before restarting the queue manager.

## **Corrective actions**

You have a choice of actions to deal with the problems of cluster error recovery. The first choice is to monitor and fix the problem, the second to monitor and postpone fixing the problem, and the final choice is to continue to manage cluster error recovery as in releases before Version 7.1.

1. Monitor the queue manager error log for the error messages AMQ9448 and AMQ5008, and fix the problem.

AMQ9448 indicates that the repository manager has returned an error after running a command. This error marks the start of trying the command again every 10 minutes, and eventually stopping the queue manager after five days, unless you postpone the shutdown.

AMQ5008 indicates that the queue manager was stopped because an IBM MQ process is missing. AMQ5008 results from the repository manager stopping after five days. If the repository manager stops, the queue manager stops.

2. Monitor the queue manager error log for the error message AMQ9448, and postpone fixing the problem.

If you disable getting messages from `SYSTEM.CLUSTER.COMMAND.QUEUE`, the repository manager stops trying to run commands, and continues indefinitely without processing any work. However, any handles that the repository manager holds to queues are released. Because the repository manager does not stop, the queue manager is not stopped after five days.

Run an MQSC command to disable getting messages from `SYSTEM.CLUSTER.COMMAND.QUEUE`:

```
ALTER QLOCAL(SYSTEM.CLUSTER.COMMAND.QUEUE) GET(DISABLED)
```

To resume receiving messages from `SYSTEM.CLUSTER.COMMAND.QUEUE` run an MQSC command:

```
ALTER QLOCAL(SYSTEM.CLUSTER.COMMAND.QUEUE) GET(ENABLED)
```

3. Revert the queue manager to the same cluster error recovery behavior as before Version 7.1.

You can set a queue manager tuning parameter to keep the queue manager running if the repository manager stops.

The tuning parameter is `TolerateRepositoryFailure`, in the `TuningParameters` stanza of the `qm.ini` file. To prevent the queue manager stopping, if the repository manager stops, set `TolerateRepositoryFailure` to `TRUE`; see Figure 67.

Restart the queue manager to enable the `TolerateRepositoryFailure` option.

If a cluster error has occurred that prevents the repository manager starting successfully, and hence the queue manager from starting, set `TolerateRepositoryFailure` to `TRUE` to start the queue manager without the repository manager.

## Special consideration

Before Version 7.1, some administrators managing queue managers that were not part of a cluster stopped the `amqrrmfa` process. Stopping `amqrrmfa` did not affect the queue manager.

Stopping `amqrrmfa` in Version 7.1 or later causes the queue manager to stop, because it is regarded as a queue manager failure. You must not stop the `amqrrmfa` process in Version 7.1 or later, unless you set the queue manager tuning parameter, `TolerateRepositoryFailure`.

## Example

```
TuningParameters:
 TolerateRepositoryFailure=TRUE
```

Figure 67. Set `TolerateRepositoryFailure` to `TRUE` in `qm.ini`

### Related information:

Queue manager configuration files, `qm.ini`

## Change in behavior of `MQS_REPORT_NOAUTH`

In IBM WebSphere MQ Version 7.1, the default behavior of `MQS_REPORT_NOAUTH` was changed to `TRUE`.

From IBM WebSphere MQ Version 7.1 onwards, this change causes UNIX platforms to behave like Windows, and log authorization failures to the error log.

Prior to IBM WebSphere MQ Version 7.1, this would only happen if you set `MQS_REPORT_NOAUTH`.

For more information, see `MQS_REPORT_NOAUTH`.

## Connect to multiple queue managers and use MQCNO\_FASTPATH\_BINDING

Applications that connect to queue managers using the MQCNO\_FASTPATH\_BINDING binding option might fail with an error and reason code MQRC\_FASTPATH\_NOT\_AVAILABLE.

An application can connect to multiple queue managers from the same process. In releases earlier than IBM WebSphere MQ Version 7.1, an application can set any one of the connections to MQCNO\_FASTPATH\_BINDING. From Version 7.1 onwards, only the first connection can be set to MQCNO\_FASTPATH\_BINDING. For the complete set of rules, see Fast path.

To assist with migration, you can set a new environment variable, AMQ\_SINGLE\_INSTALLATION. The variable reinstates the same behavior as in earlier releases, but prevents an application connecting to queue managers associated with other installations in the same process.

### Fast path

On a server with multiple installations, applications using a fast path connection to IBM WebSphere MQ Version 7.1 or later must follow these rules:

1. The queue manager must be associated with the same installation as the one from which the application loaded the IBM MQ run time libraries. The application must not use a fast path connection to a queue manager associated with a different installation. An attempt to make the connection results in an error, and reason code MQRC\_INSTALLATION\_MISMATCH.
2. Connecting non-fast path to a queue manager associated with the same installation as the one from which the application has loaded the IBM MQ run time libraries prevents the application connecting fast path, unless either of these conditions are true:
  - The application makes its first connection to a queue manager associated with the same installation a fast path connection.
  - The environment variable, AMQ\_SINGLE\_INSTALLATION is set.
3. Connecting non-fast path to a queue manager associated with a Version 7.1 or later installation, has no effect on whether an application can connect fast path.
4. You cannot combine connecting to a queue manager associated with a Version 7.0.1 installation and connecting fast path to a queue manager associated with a Version 7.1, or later installation.

With AMQ\_SINGLE\_INSTALLATION set, you can make any connection to a queue manager a fast path connection. Otherwise almost the same restrictions apply:

- The installation must be the same one from which the IBM MQ run time libraries were loaded.
- Every connection on the same process must be to the same installation. If you attempt to connect to a queue manager associated with a different installation, the connection fails with reason code MQRC\_INSTALLATION\_MISMATCH. Note that with AMQ\_SINGLE\_INSTALLATION set, this restriction applies to all connections, not only fast path connections.
- Only connect one queue manager with fast path connections.

**Related information:**

Binding options

2587 (0A1B) (RC2587): MQRC\_HMSG\_NOT\_AVAILABLE

2590 (0A1E) (RC2590): MQRC\_FASTPATH\_NOT\_AVAILABLE

**Custom scripts**

In IBM MQ for Windows, custom scripts that have been used in an earlier release to install packages might fail if any packages have been renamed, removed, or added in the later release to which you are migrating.

For an overview of server and client installation, and for links to more detailed platform-specific information about the features that are available when you are installing an IBM MQ server or client, see [Choosing what to install](#),

**Changes to data types**

A number of data types have changed between WebSphere MQ Version 7.0.1 to IBM MQ Version 8.0 and new data types have been added. This topic lists the changes for data types that have a new current version in Version 7.5.

The current version of a data type is incremented if the length of a data type is extended by adding new fields. The addition of new constants to the values that can be set in a data type does not result in a change to the current version value.

Table 94 lists the data types that have new versions. Click on the links to read about the new fields.

*Table 94. New fields added to existing data types*

| Data type                        | New version      | New fields                                                          |
|----------------------------------|------------------|---------------------------------------------------------------------|
| Channel definition               | MQCD_VERSION_10  | BatchDataLimit (MQLONG)<br>DefReconnect (MQLONG)<br>UseDLQ (MQLONG) |
| Channel exit                     | MQCXP_VERSION_8  | MCAUserSource (MQLONG)<br>pEntryPoints (PMQIEP)                     |
| Data conversion exit             | MQDXP_VERSION_2  | pEntryPoints (PMQIEP)                                               |
| Pre-connect exit                 | MQNXP_VERSION_2  | pEntryPoints (PMQIEP)                                               |
| Publish exit publication context | MQPBC_VERSION_2  | MsgDescPtr (PMQMD)                                                  |
| Publish exit                     | MQPSXP_VERSION_2 | pEntryPoints (PMQIEP)                                               |
| Cluster workload exit            | MQWXP_VERSION_4  | pEntryPoints (PMQIEP)                                               |

## Default transmission queue restriction

The product documentation in versions of IBM WebSphere MQ before Version 7.1 warned about defining the default transmission queue as `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, but no error was reported. In Version 7.1, any attempt to set or use a default transmission queue that is defined as `SYSTEM.CLUSTER.TRANSMIT.QUEUE` results in an error.

In versions of IBM WebSphere MQ before Version 7.1, no error was reported when defining the default transmission queue as `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. `MQOPEN` or `MQPUT1` MQI calls that resulted in referencing the default transmission queue did not return an error. Applications might have continued working and failed later on. The reason for the failure was hard to diagnose.

The change from Version 7.1 onwards ensures that any attempt to set the default transmission queue to `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, or use a default transmission queue set to `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, is immediately reported as an AMQ8520 error.

### Related reference:

“MQI and PCF reason code changes” on page 610

Some reason codes that affect some existing programs changed in IBM WebSphere MQ Version 7.1.

## Display channel and cluster status: Switching

From IBM WebSphere MQ Version 7.5 onwards, a cluster-sender channel that is switching its configuration to a different cluster transmission queue has a new channel state: Switching.

When migrating from a release before IBM WebSphere MQ Version 7.5 to Version 7.5 or later, existing application programs are not affected by the new state.

System management programs that monitor channel or cluster status might receive the new state as a result of an inquiry.

The state is set during the short interval while the channel modifies the destination transmission queue that messages are stored on. Before the switching state is set, messages are stored on the previously associated transmission queue. After the switching state, messages are stored on the newly configured transmission queue. The channel enters the switching state if a cluster-sender channel is starting, a configuration change is required, and the conditions for starting the switch are met.

### Related information:

Channel states

Working with cluster transmission queues and cluster-sender channels

DISPLAY CHSTATUS

Inquire Channel Status (Response)

Inquire Cluster Queue Manager (Response)

MQCHS\_\* (Command format Channel Status)

 Work with MQ Channel Status (WRKMQMCHST)

 Work with MQ Channels (WRKMQMCHL)

 Work with MQ Clusters (WRKMQMCL)

## Changes to dspmqver output

From IBM WebSphere MQ Version 7.1 onwards, new types of information are displayed by `dspmqver` to support multiple installations. The changes might affect existing administrative scripts that you have written to manage a release of IBM WebSphere MQ before IBM WebSphere MQ Version 7.1.

The changes in output from `dspmqver` that might affect existing command scripts that you have written are as follows:

1. Version 7.1 onwards has extra -f field options. If you do not specify a -f option, output from all the options is displayed. To restrict the output to the same information that was displayed in earlier releases, set the -f option to a value that was present in the earlier release. Compare the output for dspmqver in Figure 68 and Figure 69 with the output for dspmqver f 15 in Figure 70.

```
dspmqver
Name: WebSphere MQ
Version: 7.0.1.6
CMVC level: p701-L110705
BuildType: IKAP - (Production)
```

Figure 68. Default dspmqver options in IBM WebSphere MQ Version 7.0.1

```
dspmqver
Name: WebSphere MQ
Version: 7.1.0.0
Level: p000-L110624
BuildType: IKAP - (Production)
Platform: WebSphere MQ for Windows
Mode: 32-bit
O/S: Windows 7, Build 2600: SP1
InstName: 110705
InstDesc: July 5 2011
InstPath: C:\Program Files\IBM\WebSphere MQ_110705
DataPath: C:\Program Files\IBM\WebSphere MQ
Primary: No
MaxCmdLevel: 710
```

Note there are a number (1) of other installations,  
use the '-i' parameter to display them.

Figure 69. Default dspmqver options in IBM WebSphere MQ Version 7.1

```
dspmqver -f 15
Name: WebSphere MQ
Version: 7.1.0.0
Level: p000-L110624
BuildType: IKAP - (Production)
```

Figure 70. dspmqver with option to make IBM WebSphere MQ Version 7.1, or later, similar to IBM WebSphere MQ Version 7.0.1



**Related information:**

dspmqver

**Exits and installable services**

When migrating to IBM WebSphere MQ Version 7.1 or later on a distributed platform, if you install the product in a non-default location, you must update your exits and installable services. Data conversion exits generated using the **crtmqcvx** command must be regenerated using the updated command.

When writing new exits and installable services, you do not need to link to any of the following libraries:

- mqmzf
- mqm
- mqmvx
- mqmvxd
- mqic
- mqutl

For more information about updating existing exits, and writing exits and installable services that do not link to the libraries, see Writing exits and installable services on UNIX, Linux and Windows .

**Fewer IBM MQ MQI client log messages**

Before Version 7.1, an IBM WebSphere MQ MQI client used to report every failed attempt to connect to a queue manager when processing a connection name list. From Version 7.1, only if the failure occurs with the last connection in the list is a message written to the queue manager error log.

Reporting the last failure and no others, reduces growth of the queue manager error log.

**GSKit: Changes from Version 7.0 to Version 8.0**

For distributed platforms, from IBM WebSphere MQ Version 7.1, GSKit Version 8.0 is integrated with IBM WebSphere MQ and is the only version of GSKit provided with the product. GSKit Version 7.0 is no longer provided.

In IBM WebSphere MQ Version 7.0.1, if you select SSL and TLS support during installation, GSKit Version 7.0 is installed and run by default. IBM WebSphere MQ Version 7.0.1, Fix Pack 4 and later also contain an alternative, separate copy of GSKit Version 8.0 that you can install and run instead of, or in addition to, GSKit Version 7.0. From IBM WebSphere MQ Version 7.1, GSKit Version 8.0 is the only version of GSKit that is provided.

Some functions in GSKit Version 8.0 are different from the functions in GSKit Version 7.0. These differences are described in the following subtopics.

### **GSKit: Some FIPS 140-2 compliant channels do not start:**

From IBM WebSphere MQ Version 7.1 three CipherSpecs are no longer FIPS 140-2 compliant. If a client or queue manager is configured to require FIPS 140-2 compliance, channels that use the following CipherSpecs do not start after migration.

- FIPS\_WITH\_DES\_CBC\_SHA
- FIPS\_WITH\_3DES\_EDE\_CBC\_SHA
- TLS\_RSA\_WITH\_DES\_CBC\_SHA

To restart a channel, alter the channel definition to use a FIPS 140-2 compliant CipherSpec. Alternatively, configure the queue manager, or the client in the case of an IBM MQ MQI client, not to enforce FIPS 140-2 compliance.

Earlier versions of IBM WebSphere MQ enforced an older version of the FIPS 140-2 standard. The following CipherSpecs were considered FIPS 140-2 compliant in earlier versions of IBM WebSphere MQ and are also compliant in this version:

- TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA (deprecated)
- TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA
- TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (only when AltGSKit version 8 is used with Fix Pack 7.0.1.4 or later)
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA256 (only when AltGSKit version 8 is used with Fix Pack 7.0.1.4 or later)

Use these CipherSpecs if you want IBM MQ to interoperate in a FIPS 140-2 compliant manner with earlier versions.

Previous IBM MQ releases enforced an older version of the FIPS 140-2 standard. The following CipherSpecs were considered FIPS 140-2 compliant by previous IBM MQ releases and are also considered compliant in this version of IBM MQ:

- TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA (deprecated)
- TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA

Use these CipherSpecs if you need IBM MQ to interoperate in a FIPS 140-2 compliant manner with earlier IBM MQ releases.

#### **Related information:**

AMQ9196

Federal Information Processing Standards (FIPS)

Federal Information Processing Standards

FipsRequired (MQLONG)

MQSSLFIPS

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client

SSLFIPSRequired (MQLONG)

### **GSKit: Certificate Common Name (CN) not mandatory:**

In GSKit Version 8.0, the **iKeyman** command accepts any element of the distinguished name (DN), or a form of the subject alternative name (SAN). It does not mandate that you provide it with a common name. In GSKit Version 7.0, if you create a self-signed certificate using the **iKeyman** command you had to specify a common name.

The implication is that applications searching for a certificate might not be able to assume that a certificate has a common name. You might need to review how applications search for certificates, and how applications handle errors involving the common name. Alternatively, you might choose to check that all self-signed certificates are given common names.

Some other certificate tools that you might also be using, do not require a common name. It is therefore likely the change to GSKit is not going to cause you a problem.

**Related information:**

Distinguished Names

**GSKit: Commands renamed:**

The command name **gsk7cmd** is replaced with **runmqckm**; **gsk7ikm** is replaced with **strmqikm**, and **gsk7capicmd** is replaced with **runmqakm**. All the commands start the GSKit Version 8.0 certificate administration tools, and not the GSKit Version 7.0 tools.

IBM WebSphere MQ Version 7.1 or later does not use a machine-wide shared installation of GSKit: instead it uses a private GSKit installation in the IBM MQ installation directory. Each IBM WebSphere MQ Version 7.1 or later installation can use a different GSKit version. To display the version number of GSKit embedded in a particular installation, run the **dspmqrver** command from that installation as shown in the following table:

*Table 95. Renamed GSKit commands*

| Platform                | GSKit Version 7.0 command | GSKit Version 8.0 command |
|-------------------------|---------------------------|---------------------------|
| UNIX and Linux          | <b>gsk7cmd</b>            | <b>runmqckm</b>           |
| UNIX and Linux          | <b>gsk7ikm</b>            | <b>strmqikm</b>           |
| Windows, UNIX and Linux | <b>gsk7capicmd</b>        | <b>runmqakm</b>           |
| Windows, UNIX and Linux | <b>gsk7ver</b>            | <b>dspmqrver -p 64 -v</b> |

**Note:** Do not use the **gsk8ver** command to display the GSKit version number: only the **dspmqrver** command will show the correct GSKit version number.

**Related information:**

runmqckm, and runmqakm commands

Using iKeyman, iKeycmd, runmqakm, and runmqckm

dspmqrver

**GSKit: The iKeyman commands to insert a certificate do not check that all required CA certificates are present:**

The **iKeyman** command in GSKit V8.0 does not validate a certificate when it is inserted into a key repository. **iKeyman** in GSKit V7.0, validated a certificate before it inserted the certificate into a certificate store.

The implication is, that if you create a certificate using the **iKeyman** in GSKit V8.0, all the necessary intermediate and root CA certificates might not be present, or they might have expired; when the certificate is checked it might fail.

Necessary certificates might be missing, or have expired. This can cause SSL and TLS connections to fail with error AMQ9633.

### Related information:



Certificate validation and trust policy design on UNIX, Linux and Windows systems

### GSKit: PKCS#11 and JRE addressing mode:

If you use **iKeyman** or **iKeycmd** to administer certificates and keys for PKCS#11 cryptographic hardware note that the addressing mode of the JRE for these tools has changed from IBM WebSphere MQ Version 7.1.

In earlier releases of IBM WebSphere MQ, on the following platforms the JRE was 32 bit however in IBM WebSphere MQ Version 7.1 or later, it is 64 bit only. Where it has changed you might need to install additional PKCS#11 drivers appropriate for the addressing mode of the **iKeyman** and **iKeycmd** JRE. This is because the PKCS#11 driver must use the same addressing mode as the JRE. The following table shows the IBM WebSphere MQ Version 7.1 or later JRE addressing modes.

Table 96. IBM WebSphere MQ Version 7.1 or later JRE addressing modes

| Platform                                                                                  | JRE Addressing Mode |
|-------------------------------------------------------------------------------------------|---------------------|
| Windows (32 bit or 64 bit)                                                                | 32                  |
| Linux for System x 32 bit                                                                 | 32                  |
| Linux for System x 64 bit                                                                 | 64                  |
| Linux on POWER Systems - Big Endian                                                       | 64                  |
| Linux for IBM Z                                                                           | 64                  |
| HP-UX                                                                                     | 64                  |
| Solaris Sparc                                                                             | 64                  |
| Solaris x86-64                                                                            | 64                  |
| AIX                                                                                       | 64                  |
|  z/OS  | n/a                 |
|  IBM i | n/a                 |

### GSKit: Import of a duplicate PKCS#12 certificate:

In GSKit V8.0, the **iKeyman** command does not report an attempt to import a duplicate PKCS#12 certificate as an error. In GSKit V7.0, the **iKeyman** command reported an error. In neither version is a duplicate certificate imported.

For GSKIT V8.0, a duplicate certificate is a certificate with the same label and public key.

The implication is that if some of the issuer information is different, but the name and public key are the same, the changes are not imported. The correct way to update a certificate is to use the **-cert -receive** option, which replaces an existing certificate.

**gskcapicmd** does not allow or ignore duplicates on import in this way.

**Related information:**

Importing a personal certificate into a key repository on UNIX, Linux or Windows systems

**GSKit: Certificate stores created by iKeyman and iKeycmd no longer contain CA certificates:**

The **iKeyman** and **iKeycmd** utilities in GSKit V8.0 create a certificate store without adding pre-defined CA certificates to the store. To create a working certificate store, you must now add all the certificates that you require and trust. In GSKit V7.0 **iKeyman** and **iKeycmd** created a certificate store that already contained CA certificates.

Existing data bases created by GSKit V7.0 are unaffected by this change.

**Related information:**

Adding default CA certificates into an empty key repository, on UNIX, Linux or Windows systems with GSKit Version 8.0

**GSKit: Password expiry to key database deprecated:**

In GSKit V8.0, the password expiry function in **iKeyman** continues to work the same as in GSKit V7.0, but it might be withdrawn in future versions of GSKit.

Use the file system protection provided with the operating system to protect the key database and password stash file.

**Linux: Recompile C++ applications and update run time libraries:**

C++ IBM MQ MQI client and server applications on Linux must be recompiled using a supported version of the GNU Compiler Collection (GCC). The C++ run time libraries for this version of GCC must be installed in `/usr/lib` or `/usr/lib64`.

For information about which versions of GCC are supported , see IBM MQ System Requirements, and follow the links for Linux. Older versions of the GCC that are not included in the System Requirements information are no longer supported.

If you are using one of the supported Linux distributions, the libraries are correctly installed; see IBM MQ System Requirements.

The GCC Version 4.x libraries support SSL and TLS connections from an IBM MQ MQI client. SSL and TLS use GSKit Version 8.0, which depends on `libstdc++.so.6`. `libstdc++.so.6` is included in GCC 4.x.

### **GSKit: Signature algorithm moved out of settings file:**

In GSKit V8.0, the default signature algorithm used when creating self-signed certificates or certificate requests or selected in the creation dialogs is passed as a command-line parameter. In GSKit V7.0, the default signature algorithm was specified in the settings file.

The change has very little effect: it causes a different default signature algorithm to be selected. It does not alter the selection of a signature algorithm.

#### **Related information:**

`runmqckm` and `runmqakm` options

Creating a self-signed personal certificate on UNIX, Linux, and Windows systems

### **GSKit: Signed certificate validity period not within signer validity:**

In GSKit V8.0, the `iKeyman` command does not check whether the validity period of a resulting certificate is within the validity period of the signed certificate. In GSKit V7.0, `iKeyman` checked that the validity period of the resulting certificate was within the validity period of the signed certificate.

The IETF RFC standards for SSL/TLS allow a certificate whose validity dates extend beyond those of its signer. This change to GSKit brings it into line with those standards. The check is whether the certificate is issued within the validity period of the signer, and not whether it expires within the validity period of the signer.

#### **Related information:**

How SSL and TLS provide identification, authentication, confidentiality, and integrity

### **GSKit: Stricter default file permissions:**

The default file permissions set by `runmqckm` and `strmqikm` in IBM WebSphere MQ Version 7.5 on UNIX and Linux are stricter than the permissions that are set by `runmqckm`, `strmqikm`, `gsk7cmd`, and `gsk7ikm` in earlier releases of IBM MQ.

The permissions set by `runmqckm` and `strmqikm` in IBM WebSphere MQ Version 7.5 permit only the creator to access the UNIX and Linux SSL/TLS key databases. The `runmqckm`, `strmqikm`, `gsk7cmd`, and `gsk7ikm` tools in earlier releases of IBM MQ set world-readable permissions, making the files liable to theft and impersonation attacks.

The permissions set by `gsk7capicmd`, in earlier releases of IBM MQ, and `runmqakm` in IBM WebSphere MQ Version 7.5, permit only the creator to access UNIX and Linux SSL/TLS key databases.

The migration of SSL/TLS key databases to Version 7.5 does not alter their access permissions. In many cases, administrators set more restrictive access permissions on these files to overcome the liability to theft and impersonation attacks; these permissions are retained.

The default file permissions set on Windows are unchanged. Continue to tighten up the access permissions on SSL/TLS key database files on Windows after creating the files with `runmqckm` or `strmqikm`.

**Related information:**

Accessing and securing your key database files on Windows

Accessing and securing your key database files on UNIX and Linux systems

**Java: Change in behavior of default value of MQEnvironment.userID**

Change caused when using CLIENT transport for a channel not have a security exit defined.

From IBM WebSphere MQ Version 7.1, if an IBM WebSphere MQ classes for Java application is connecting to a queue manager, using the CLIENT transport through a channel that does not have a security exit defined, and the MQEnvironment.userID field is kept at its default value of the empty string (""), the IBM WebSphere MQ classes for Java application queries the value of the Java System Property user.name and passes this to the queue manager for authorization as part of the MQQueueManager constructor.

If the user specified by the Java System Property user.name is not authorized to access the queue manager, the MQQueueManager constructor throws an MQException containing Reason Code MQRC\_NOT\_AUTHORIZED.

**Java: Different message property data type returned**

From IBM WebSphere MQ Version 7.1, if the data type of a message property is set, the same data type is returned when the message is received. This is different from IBM WebSphere MQ Version 7.0.1, where in some circumstances properties set with a specific type were returned with the default type String.

The change affects Java applications that used the MQRFH2 class, and retrieved properties using the getFieldValue method.

You can write a message property in Java using a method such as setIntFieldValue. In IBM WebSphere MQ Version 7.0.1 the property is written into an MQRFH2 header with a default type of String. When you retrieve the property with the getFieldValue method, a String object is returned.

The change is that now the correct type of object is returned, in this example the type of object returned is Integer.

If your application retrieves the property with the getIntFieldValue method, there is no change in behavior; an int is returned. If property is written to the MQRFH2 header by some other means, and the data type is set, then getFieldValue returns the correct type of object.

**Related information:**

Class MQRFH2

**JMS: Change in behavior of the default user identifier value**

Change caused when using CLIENT transport for a channel that does not have a security exit defined.

From IBM WebSphere MQ Version 7.1, if an IBM WebSphere MQ classes for JMS application is connecting to a queue manager, using the CLIENT transport through a channel that does not have a security exit defined, and no user identifier is specified, by calling ConnectionFactory.createConnection(), the IBM WebSphere MQ classes for JMS application queries the value of the Java System Property user.name and passes this to the queue manager for authorization as part of the call to create a connection from a connection factory object. This behavior also occurs when calling ConnectionFactory.createConnection(String, String) and passing a blank or null value for the first parameter **userID**.

If the user specified by the Java System Property user.name is not authorized to access the queue manager, a JMSEException containing Reason Code MQRC\_NOT\_AUTHORIZED is thrown.

## **JMS: some objects are no longer serializable**

JMS objects such as JMS Connections and JMS Sessions, which used to be serializable when using the IBM MQ classes for JMS for IBM WebSphere MQ Version 7.0.1, are no longer serializable when using IBM WebSphere MQ Version 7.1 or later.

When a Java application serializes an object, state information about that object is written to an output stream, such as a file. The contents of output stream can then be read at a later date, to reconstruct (or deserialize) the Java object so that it can be reused.

The following interfaces provided by the IBM MQ classes for JMS are implemented by objects that represent an active connection from an application to an IBM MQ queue manager:

- JMSConnection
- JMSQueueConnection
- JMSTopicConnection
- JMSSession
- JMSQueueSession
- JMSTopicSession

There is an IBM MQ connection handle (hconn) associated with every JMSConnection, JMSQueueConnection, JMSTopicConnection, JMSSession, JMSQueueSession and JMSTopicSession object that is created by an application.

Serializing one of these objects would result in state information about those objects being written to an output stream. This included information about the connection handle to IBM MQ that was associated with the object.

However, there was no guarantee that the connection handle would still be valid when the object was deserialized and reused, which could lead to unexpected behavior. To prevent applications running into these issues, the IBM MQ classes for JMS for Version 7.1 (and later) will throw a `NotSerializableException` if an application tries to:

- Serialize a JMSConnection, JMSQueueConnection, JMSTopicConnection, JMSSession, JMSQueueSession or JMSTopicSession object using the method:  
`writeObject(ObjectOutputStream)`
- Deserialize a JMSConnection, JMSQueueConnection, JMSTopicConnection, JMSSession, JMSQueueSession or JMSTopicSession object using the method:  
`readObject(ObjectInputStream)`



## JMS: Reason code changes

From IBM WebSphere MQ Version 7.1, some reason codes returned in JMS exceptions have changed. The changes affect MQRC\_Q\_MGR\_NOT\_AVAILABLE and MQRC\_SSL\_INITIALIZATION\_ERROR.

Before IBM WebSphere MQ Version 7.1, if a JMS application call fails to connect, it receives an exception with a reason code 2059 (080B) (RC2059): MQRC\_Q\_MGR\_NOT\_AVAILABLE. From Version 7.1, the application can still receive MQRC\_Q\_MGR\_NOT\_AVAILABLE, or one of the following more specific reason codes.

- 2537 (09E9) (RC2537): MQRC\_CHANNEL\_NOT\_AVAILABLE
- 2538 (09EA) (RC2538): MQRC\_HOST\_NOT\_AVAILABLE
- 2539 (09EB) (RC2539): MQRC\_CHANNEL\_CONFIG\_ERROR
- 2540 (09EC) (RC2540): MQRC\_UNKNOWN\_CHANNEL\_NAME

Similarly, when trying to connect, a JMS application might have received 2393 (0959) (RC2393): MQRC\_SSL\_INITIALIZATION\_ERROR. From Version 7.1, the application can still receive MQRC\_SSL\_INITIALIZATION\_ERROR, or a more specific reason code, such as 2400 (0960) (RC2400): MQRC\_UNSUPPORTED\_CIPHER\_SUITE, that identifies the cause of the SSL initialization error.

## JMS: ResourceAdapter object configuration

When WebSphere Application Server connects to IBM MQ it creates message driven beans (MDBs) using JMS connections. From Version 7.1, these MDBs can no longer share one JMS connection. The configuration of ResourceAdapter object is migrated so that there is a single MDB for each JMS connection.

## Changed ResourceAdapter properties

### connectionConcurrency

The maximum number of MDBs to share a JMS connection. Sharing connections is not possible and this property always has the value 1. Its previous default value was 5.

### maxConnections

This property is the number of JMS connections that the resource adapter can manage. From Version 7.1, it also determines the number of MDBs that can connect because each MDB requires one JMS connection. The default value of maxConnections is now 50. Its previous default value was 10.

If connectionConcurrency is set to a value greater than 1, the maximum number of connections supported by the resource adapter is scaled by the value of connectionConcurrency. For example, if maxConnections is set to 2 and connectionConcurrency is set to 4, the maximum number of connections supported by the resource adapter is 8. As a result, connectionConcurrency is set to 1 and maxConnections is set to 8.

If connectionConcurrency is set to a value greater than 1, it is adjusted automatically. To avoid automatic adjustment, set connectionConcurrency to 1. You can then set maxConnections to the value you want.

The scaling mechanism ensures that sufficient connections are available for existing deployments whether you have changed them in your deployment, configuration, or programs.

If the adjusted maxConnections value exceeds the MAXINST or MAXINSTC attributes of any used channel, previously working deployments might fail.

The default value of both channel attributes equates to unlimited. If you changed them from the default value, you must ensure that the new maxConnections value does not exceed MAXINST or MAXINSTC.

**Related information:**

Configuration of the ResourceAdapter object

**V8.0.0.4** Maximum instances (MAXINST)

Maximum instances per client (MAXINSTC)

**MQI and PCF reason code changes**

Some reason codes that affect some existing programs changed in IBM WebSphere MQ Version 7.1.

**MQRC\_NOT\_OPEN\_FOR\_INPUT**

In IBM WebSphere MQ Version 7.0 a queue opened with `MQOO_OUTPUT`, and then browsed, returned an error with the wrong reason-code, `MQRC_NOT_OPEN_FOR_INPUT`. The correct reason-code, `MQRC_NOT_OPEN_FOR_BROWSE`, was issued by Version 6.0 and earlier. Version 7.1 and later correctly returns an error with the same reason code as Version 6.0, `MQRC_NOT_OPEN_FOR_BROWSE`.

**MQRC\_DEF\_XMIT\_Q\_USAGE\_ERROR**

The product documentation in versions of IBM WebSphere MQ before Version 7.1 warned about defining the default transmission queue as `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, but no error was reported. From Version 7.1 onwards, an attempt to open the default transmission queue, defined as `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, results in the error `MQRC_DEF_XMIT_Q_USAGE_ERROR`.

**MQRC\_FASTPATH\_NOT\_AVAILABLE**

An application that connects to multiple queue managers in the same process and uses `MQCNO_FASTPATH_BINDING` might fail with an error and reason code `MQRC_FASTPATH_NOT_AVAILABLE`; see “Connect to multiple queue managers and use `MQCNO_FASTPATH_BINDING`” on page 597.

**MQRCCF\_DEF\_XMIT\_Q\_CLUS\_ERROR**

The product documentation in versions of IBM WebSphere MQ before Version 7.1 warned about defining the default transmission queue as `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, but no error was reported. From Version 7.1 onwards, an attempt to alter the queue manager attribute `DEFXMITQ` to `SYSTEM.CLUSTER.TRANSMIT.QUEUE` results in an error. The PCF reason code is 3269 (0CC5) (RC3269): `MQRCCF_DEF_XMIT_Q_CLUS_ERROR`.

**Related reference:**

“Connect to multiple queue managers and use `MQCNO_FASTPATH_BINDING`” on page 597

Applications that connect to queue managers using the `MQCNO_FASTPATH_BINDING` binding option might fail with an error and reason code `MQRC_FASTPATH_NOT_AVAILABLE`.

**Related information:**

2036 (07F4) (RC2036): `MQRC_NOT_OPEN_FOR_BROWSE`

2037 (07F5) (RC2037): `MQRC_NOT_OPEN_FOR_INPUT`

2590 (0A1E) (RC2590): `MQRC_FASTPATH_NOT_AVAILABLE`

3269 (0CC5) (RC3269): `MQRCCF_DEF_XMIT_Q_CLUS_ERROR`

AMQ8520

**Publish/Subscribe: Delete temporary dynamic queue**

If a subscription is associated with a temporary dynamic queue, when the queue is deleted, the subscription is deleted. Publish/subscribe applications migrated from WebSphere Message Broker are unchanged. The change does not affect the behavior of integrated publish/subscribe applications, which are written using the MQI publish/subscribe interface.

- In IBM WebSphere MQ Version 7.0, if you migrate or create a queued publish/subscribe application that uses `MQRFH1`, it behaves the same as IBM WebSphere MQ Version 6.0. You can create a temporary dynamic queue for a subscription, and if the queue is deleted, the subscription is not deleted, as in IBM WebSphere MQ Version 6.0. The lack of a subscriber queue results in any matching publications ending up on the dead letter queue.

- From IBM WebSphere MQ Version 7.0.1, Fix Pack 6 onwards, in the same MQRFH1 queued publish/subscribe case, if the temporary dynamic queue is deleted, the subscription is deleted. This change prevents a buildup of publications from a subscription without a subscriber queue ending up on the dead letter queue.
- From IBM WebSphere MQ Version 7.0 onwards, if you migrate or create a queued publish/subscribe application that uses MQRFH2, it behaves the same as WebSphere Message Broker IBM Integration Bus. You can create a temporary dynamic queue for a subscription, and if the queue is deleted, the subscription is deleted, as in IBM Integration Bus. MQRFH2 publish/subscribe applications are typically migrated from IBM Integration Bus.
- From IBM WebSphere MQ Version 7.0 onwards, if you create a durable subscription using integrated publish/subscribe, you cannot define a temporary dynamic queue as the destination for its matching publications.
- From IBM WebSphere MQ Version 7.0, you can create a managed, non-durable subscription using integrated publish/subscribe, which creates a temporary dynamic queue as the destination for matching publications. The subscription is deleted with the queue.

## Summary

From IBM WebSphere MQ Version 7.0.1 onwards, you cannot create a temporary dynamic queue as the destination for publications for a durable subscription using the integrated publish/subscribe interface.

If you use either of the queued publish/subscribe interfaces, MQRFH1 or MQRFH2, the behavior is the same. You can create a temporary dynamic queue as the subscriber queue, and if the queue is deleted, the subscription is deleted with it.

## Queue manager logs: Default sizes increased

From IBM WebSphere MQ Version 7.1, the default size of a queue manager log files changed to 4096. The queue manager error log increased from 256 KB to 2 MB on some platforms. The change affects both new and migrated queue managers.

### Queue manager log

From IBM WebSphere MQ Version 7.1, the default queue manager log size is 4096. For more information on setting non-default values, see The IBM MQ configuration file, `mqs.ini`.

### Queue manager error log

From IBM WebSphere MQ Version 7.1, the `AMQERR nn.log` increased from 256 KB to 2 MB on


 IBM i, UNIX, Linux, and Windows platforms.

You can override the change either by setting the environment variable `MQMAXERRORLOGSIZE`, or by setting `ErrorLogSize` in the `QMErrorLog` stanza in the `qm.ini` file.

The change increases the number of error messages that are saved in the error logs.

## Related information:

Queue manager error logs

 The queue manager error log stanza

## Removal of dspmqsver command

Before IBM WebSphere MQ Version 7.5, the **dspmqsver** command was used to display the version of IBM WebSphere MQ Advanced Message Security.

From Version 7.5, IBM WebSphere MQ Advanced Message Security is a separately licensed component of the main product and the version information is displayed as part of the **dspmqver** command.

## Security: SSLPEER and SSLCERTI changes

IBM WebSphere MQ Version 7.1 or later obtains the Distinguished Encoding Rules (DER) encoding of the certificate and uses it to determine the subject and issuer distinguished names. The subject and issuer distinguished names are used in the SSLPEER and SSLCERTI fields. A SERIALNUMBER attribute is also included in the subject distinguished name and contains the serial number for the certificate of the remote partner. Some attributes of subject and issuer distinguished names are returned in a different sequence from releases before Version 7.1.

The change to subject and issuer distinguished names affects channel security exits. It also affects applications which depend upon the subject and issuer distinguished names that are returned by the PCF programming interface. Channel security exits and applications that set or query SSLPEER and SSLCERTI must be examined, and possibly changed. The fields that are affected are listed in Table 97 and Table 98.

*Table 97. Channel status fields affected by changes to subject and issuer distinguished names*

| Channel status attribute | PCF channel parameter type  |
|--------------------------|-----------------------------|
| SSL Peer (SSLPEER)       | MQCACH_SSL_SHORT_PEER_NAME  |
| SSLCERTI                 | MQCACH_SSL_CERT_ISSUER_NAME |

*Table 98. Channel data structures affected by changes to subject and issuer distinguished names*

| Channel data structure         | Field                          |
|--------------------------------|--------------------------------|
| MQCD - Channel definition      | SSLPeerNamePtr (MQPTR)         |
| MQCXP - Channel exit parameter | SSLRemCertIssNamePtr (PMQVOID) |

Existing peer name filters specified in the SSLPEER field of a channel definition are not affected. They continue to operate in the same manner as in earlier releases. The peer name matching algorithm has been updated to process existing SSLPEER filters. It is not necessary to alter any channel definitions.

**Related information:**

Channel security exit programs

**Security: Disable IBM MQ AMS at the client**

From IBM WebSphere MQ Version 7.5, you can disable IBM WebSphere MQ Advanced Message Security at the client, for Java and C clients, to prevent errors when they are connecting to queue managers that are running on earlier versions of IBM WebSphere MQ.

For Java clients, in IBM WebSphere MQ Version 7.5.0.4 and later, you can disable IBM WebSphere MQ Advanced Message Security by setting an environment variable `AMQ_DISABLE_CLIENT_AMS`.

For C clients, in IBM WebSphere MQ Version 7.5.0.5 and later, you can disable IBM WebSphere MQ Advanced Message Security by using the `DisableClientAMS` property, under the **Security** stanza in the `mqclient.ini` file.

**Related information:**

Environment variables used to disable IBM MQ AMS at the client


**Telemetry: Installer integrated with IBM MQ**

From Version 7.1, IBM MQ Telemetry is no longer installed separately from IBM MQ. It is installed as a component of the main product. If you installed IBM MQ Telemetry with Version 7.0.1, you must uninstall it before installing Version 7.1 or later.

You can install IBM MQ Telemetry at the same time as IBM MQ, or you can rerun the installer and install IBM MQ Telemetry at a later time.

**Related information:**

Installing IBM MQ Telemetry

 [Uninstalling IBM MQ Telemetry Version 7.0.1 components](#)

**Telemetry: Support for the MQTT protocol over WebSockets**

IBM WebSphere MQ Version 7.5.0, Fix Pack 1 and later supports the MQTT protocol over WebSockets. This enables it to be a server for clients using the MQTT messaging client for JavaScript.

A new communication protocol parameter (`PROTOCOL`) has been added to the MQTT channel definition (`DEFINE CHANNEL (MQTT)`):

- If the parameter is set to `MQTTV3`, the channel only accepts connections from clients using Version 3 of the MQ Telemetry Transfer protocol. This was the only protocol supported before IBM WebSphere MQ Version 7.5.0, Fix Pack 1.
- If the parameter is set to `HTTP`, the channel only accepts HTTP requests for pages, or WebSockets connections to IBM MQ Telemetry.
- If the parameter is set to `MQTTV3,HTTP`, the channel accepts connections from clients using either protocol. This is the default behavior for new MQTT channels created with IBM WebSphere MQ Version 7.5.0, Fix Pack 1 and later versions.

When a client connects to an MQTT channel using SSL, the parameter `SSLCAUTH` determines whether IBM MQ requires a certificate from the client (see `DEFINE CHANNEL (MQTT)` ). Before IBM WebSphere MQ Version 7.5.0, Fix Pack 1, this parameter could be either `REQUIRED` or `OPTIONAL` for MQTT channels:

- `REQUIRED` means that IBM MQ requests a certificate from the client and the client must supply a valid certificate.
- `OPTIONAL` means that IBM MQ will request a certificate from the client but the client does not have to supply one. The client connection is allowed if the client supplies a valid certificate or if the client does not supply a certificate. The client connection is disallowed only if the client supplies an invalid certificate.

In IBM WebSphere MQ Version 7.5.0, Fix Pack 1 and later, the parameter SSLCAUTH can be set to NEVER for MQTT channels. NEVER means that IBM MQ never requests a certificate from the client. The new value was added as part of the support for clients using the MQTT messaging client for JavaScript. It accommodates the behavior of some web browsers which treat the request for a client certificate as a protocol error.

**Related information:**

Getting started with the MQTT messaging client for JavaScript

## **MQ Explorer changes**

From Version 7.1, IBM WebSphere Eclipse Platform it is not required to run MQ Explorer and is therefore no longer shipped with the product. The change makes no difference to administrators who run MQ Explorer. For developers who run MQ Explorer in an Eclipse development environment, a change is necessary. You must install and configure a separate Eclipse environment to be able to switch between MQ Explorer and other perspectives.

## **Packaging changes**

In versions of IBM WebSphere MQ earlier than Version 7.1, you can select the Workbench mode preference in MQ Explorer. In workbench mode, you could switch to the other perspectives installed in the WebSphere Eclipse Platform. In Version 7.1 or later, you can no longer set the Workbench mode preference, because the WebSphere Eclipse Platform is not shipped with MQ Explorer.

To switch between MQ Explorer and other perspectives, you must install MQ Explorer into your own Eclipse environment or into an Eclipse-based product. You can then switch between perspectives. For example, you can develop applications using IBM MQ classes for JMS or IBM MQ Telemetry applications; see [Creating your first IBM MQ Telemetry Transport publisher application using Java](#)

If you installed extensions to previous versions of MQ Explorer, such as SupportPacs or WebSphere Message Broker (now known as IBM Integration Bus ) Explorer, you must reinstall compatible versions of the extensions after upgrading MQ Explorer to Version 7.1 or later.

If you continue to run IBM WebSphere MQ Version 7.0.1 on the same server as a later version of the product, and you use MQ Explorer, each installation uses its own installation of MQ Explorer. When you uninstall Version 7.0.1, its version of MQ Explorer is uninstalled. To remove IBM WebSphere Eclipse Platform, uninstall it separately. The workspace is not deleted.

## **Test result migration**

Test results are not migrated from version to version. To view any test results, you must rerun the tests.

**Related information:**

Multiple installations of MQ Explorer

## **AIX: Shared objects**

On AIX, for Version 7.1. and later, the .a shared objects in the lib64 directory contains both the 32 bit and 64 bit objects. A symlink to the .a file is also placed in the lib directory. The AIX loader can then correctly pick up the correct object for the type of application being run.

This means that IBM MQ applications can run with the LIBPATH containing either the lib or lib64 directory, or both.

## **AIX: /usr/lpp/mqm symbolic link removed**

Before Version 6.0, IBM WebSphere MQ placed a symbolic link in /usr/lpp/mqm on AIX. The link ensured queue managers and applications migrated from IBM WebSphere MQ versions before Version 5.3 continued to work, without change. However, this link is not created in Version 7.1, or later.

In version 5.0, IBM MQ for AIX was installed into /usr/lpp/mqm. That changed in Version 5.3 to /usr/mqm. A symbolic link was placed in /usr/lpp/mqm, linking to /usr/mqm. Existing programs and scripts that relied on the installation into /usr/lpp/mqm continued to work unchanged. That symbolic link was removed in Version 7.5, because you can now install IBM MQ in any directory. Applications and command scripts are affected by the change.

The effect on applications is no different to the effect of migrating on other UNIX and Linux platforms. If the installation is made primary, then symbolic links to the IBM MQ link libraries are placed in /usr/lib. Most applications migrated from earlier IBM MQ versions search the default search path, which normally includes /usr/lib. The applications find the symbolic link to the IBM MQ load libraries in /usr/lib.

If the installation is not primary, then you must configure the correct search path to load the IBM MQ link libraries. If you choose to run **setmqenv**, IBM MQ places the IBM MQ link library path into LIBPATH. Unless the application is configured not to search the LIBPATH, if for example it is a setuid or setgid application, then the IBM MQ library is loaded successfully; see “UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542.

If you have written command scripts that run IBM MQ commands, you might have coded explicit paths to the directory tree where IBM MQ was installed. You must modify these command scripts. You can run **setmqenv** to create the correct environment to run the command scripts. If you have set the installation as primary, you do not have to specify the path to the command.

### **Related tasks:**

“UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542  
Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

### **Related information:**

“Loading IBM MQ libraries” on page 461

## **AIX, HP-UX, and Solaris: Building applications for TXSeries**

From Version 7.1, applications that link to TXSeries® must load the mqzi\_r library, not the mqz\_r library. These applications must load mqzi\_r from the installation that is associated with the queue manager to which the application is connected.

Before Version 7.1, IBM WebSphere MQ applications that used the TXSeries CICS support loaded the IBM WebSphere MQ library, mqz\_r. From Version 7.1 those applications must load the IBM WebSphere MQ library, mqzi\_r instead. You must change your build scripts accordingly and rebuild your applications.

In Version 7.1 and later, mqz\_r includes code to load a different version of the library. IBM MQ loads a different version of the library, if it detects that the queue manager to which the application is connected is associated with a different installation from the installation from which the library was loaded. However, mqzi\_r does not include the additional code. When using TXSeries, the application must run with the library it that loaded, and not a different library loaded by IBM MQ. For this reason, applications that use the TXSeries CICS support must load the mqzi\_r library, and not the mqz\_r library.

An implication of applications loading mqzi\_r instead of mqz\_r is that an application must load the correct version of mqzi\_r. The application must load the version of mqzi\_r from the installation that is associated with the queue manager to which the application is connected.

**Related information:**

Building libraries for use with TXSeries for Multiplatforms version 5

AIX: TXSeries CICS support

HP-UX: TXSeries CICS support

Solaris: TXSeries CICS support

**Linux: Recompile C++ applications and update run time libraries**

C++ IBM MQ MQI client and server applications on Linux must be recompiled using a supported version of the GNU Compiler Collection (GCC). The C++ run time libraries for this version of GCC must be installed in `/usr/lib` or `/usr/lib64`.

For information about which versions of GCC are supported, see IBM MQ System Requirements, and follow the links for Linux. Older versions of the GCC that are not included in the System Requirements information are no longer supported.

If you are using one of the supported Linux distributions, the libraries are correctly installed; see IBM MQ System Requirements.

The GCC Version 4.x libraries support SSL and TLS connections from an IBM MQ MQI client. SSL and TLS use GSKit Version 8.0, which depends on `libstdc++.so.6`. `libstdc++.so.6` is included in GCC 4.x.

**Linux: Increased shared memory allocation required**

The maximum amount of shared memory (SHMMAX) to allocate on Linux systems. The default system allocation is 32 MB. IBM MQ starts by allocating 64 MB and increases its allocation on demand by doubling its previous allocation. On a production system set SHMMAX to at least 256 MB to accommodate additional allocations.

**Related information:**

Additional settings for installing IBM MQ on Linux systems

**UNIX and Linux : `crtmq1nk` and `d1tmq1nk` removed**

Before Version 7.1, the `crtmq1nk` and `d1tmq1nk` commands created symbolic links in subdirectories of `/usr`. From Version 7.1 onwards, you must use the `setmqinst` command instead.

**Related reference:**

“UNIX and Linux: `/usr` symbolic links removed” on page 618

From Version 7.1, on all UNIX and Linux platforms, the links from the `/usr` file system are no longer made automatically. In order to take advantage of these links, you must set an installation as the primary installation.

**Related information:**

Changing the primary installation

`setmqinst`

**UNIX and Linux: Message catalogs moved**

From Version 7.1, message catalogs are no longer stored in the system directories. To support multiple installations, copies of the message catalogs are stored with each installation. If you are migrating from a release before Version 7.1 and you want messages only in the locale of your system, the change has no affect on your system. If you have customized the way in which the search procedure selects a message catalog, then the customization might no longer work correctly.

Set the LANG environment variable to load a message catalog for a different language from the system locale.



### Related information:

Displaying messages in your national language on UNIX and Linux systems

### UNIX and Linux: MQ services and triggered applications

From Version 7.1, both `LD_LIBRARY_PATH` and `$ORIGIN` work for MQ services and triggered applications. For this reason, for Version 7.1 or later, MQ Services and triggered applications run under the user ID that started the queue manager and not `setuid` or `setgid`.

If you are migrating from a release before Version 7.1, if any files used by the service were previously restricted to certain users, then they might not be accessible by the user ID that started the queue manager. Resources used by MQ services or triggered applications must be adjusted as appropriate.

#### Note:

- On AIX, `LD_LIBRARY_PATH` is also known as `LIBPATH` and `$ORIGIN` is not supported.
- On HP-UX, `LD_LIBRARY_PATH` is also known as `SHLIB_PATH`.

### Related information:

Working with services

Starting IBM MQ applications using triggers

### UNIX and Linux: `ps -ef | grep amq` interpretation

From Version 7.1, the interpretation of the list of IBM MQ processes that results from filtering a scan of UNIX or Linux processes has changed. The results can show IBM MQ processes running for multiple installations on a server. Before Version 7.1, the search identified IBM MQ processes running on only a single installation of the product on a UNIX or Linux server.

The implications of this change depend on how the results are qualified and interpreted, and how the list of processes is used. The change affects you, only if you start to run multiple installations on a single server. If you have incorporated the list of IBM MQ processes into administrative scripts or manual procedures, you must review the usage.

### Examples

The following two examples, which are drawn from the product documentation, illustrate the point.

1. In the product documentation, before Version 7.1, the scan was used as a step in tasks to change the installation of IBM MQ. The purpose was to detect when all queue managers had ended. In Version 7.1 or later, the tasks use the `dspmq` command to detect when all queue managers associated with a specific installation have ended.
2. In the product documentation, a process scan is used to monitor starting a queue manager in a high availability cluster. Another script is used to stop a queue manager. In the script to stop a queue manager, if the queue manager does not end within a period of time, the list of processes is piped into a `kill -9` command. In both these cases, the scan filters on the queue manager name, and is unaffected by the change to multiple installations.

**Related information:**

Stopping a queue manager under the control of a high availability (HA) cluster  
 Monitoring a queue manager

**UNIX and Linux: /usr symbolic links removed**

From Version 7.1, on all UNIX and Linux platforms, the links from the /usr file system are no longer made automatically. In order to take advantage of these links, you must set an installation as the primary installation.

In previous releases the installation of the product on UNIX and Linux created the symbolic links shown in Table 99. In Version 7.1 or later, these links are not created. You must run **setmqinst** to create a primary installation containing symbolic links. No symbolic links are created in other installations.

*Table 99. Default symbolic links created in releases before Version 7.1*

| Symbolic link from  | To                  |
|---------------------|---------------------|
| /usr/bin/amq...     | /opt/mqm/bin/amq... |
| /usr/lib/amq...     | /opt/mqm/lib/amq... |
| /usr/include/cmq... | /opt/mqm/inc/cmq... |
| /usr/share/man/...  | /opt/mqm/man/...    |

Only a subset of those links created with previous releases are now made, see External library and control command links to primary installation on UNIX and Linux

For more information about choosing whether to have a primary installation, see Choosing a primary installation

**Related tasks:**

“UNIX: Migrating IBM MQ library loading from Version 7.0.1, or later, to the latest version” on page 542  
 Investigate whether applications connecting to the latest version of the product are linked to, and load libraries from, the correct installation.

**Related reference:**

Primary installation  
 Explanation of the IBM MQ primary installation.

**Related information:**

Changing the primary installation  
 setmqinst

**Windows and UNIX: Configurable certificate validation policy**

From, Version 7.1.0, Fix Pack 2 you configure the product to specify which SSL or TLS certificate validation policy is used to validate digital certificates received from remote partner systems.

**Configurable certificate validation policy on Windows and UNIX platforms**

If you need to ensure that your certificate validation is compliant with RFC 5280 for improved security in certificate validation, see Configuring certificate validation policies in IBM MQ.

See Certificate validation policies in IBM MQ for more information about IBM MQ certificate validation policies.

## Windows: amqmsrvn.exe process removed

From Version 7.1, the amqmsrvn.exe DCOM process was replaced by a Windows service, amqsvc.exe. This change is unlikely to cause any problems. However, you might have to make some changes. You might have configured the user that runs the Windows service MQSeriesServices without the user right to “Log on as a service”. Alternatively, the user might not have “List Folder” privilege on all the subdirectories from the root of the drive to the location of the service amqsvc.exe.

If you omitted the “Log on as a service” user privilege, or one of the subdirectories under which the product is installed does not grant the “List Folder” privilege to the user, the MQ\_InstallationName IBM MQ Windows services in Version 8.0 fails to start.

## Diagnosing the problem

If the service fails to start, Windows event messages are generated:

- If you did not give the user the “Log on as a service” user privilege, the Windows Service Control Manager adds an event: 7038: The user has not been granted the requested logon type . The **strmqsvc** command reports error 1069.
- If you did not give the user the “List Folder” privilege, the Windows Service Control Manager adds an event: 7009: Timed out waiting for the service to connect . The **strmqsvc** command reports error 1053.

If the “Prepare IBM MQ” wizard encounters a failure when validating the security credentials of the user performing an installation, an error is returned: IBM WebSphere MQ is not correctly configured for Windows domain users . This error indicates that the service failed to start.

## Resolution

To resolve this problem:

- Check that the user has the “Log on as a service” user privilege
- Check that the user is a member of the local mqm group
- Check that the local mqm group has “List Folder” privilege on each subdirectory in the path to the service amqsvc.exe.

### Related reference:

“Windows: “Logon as a service” required” on page 620

From Version 7.1, the user ID that runs the IBM MQ Windows service must have the user privilege to “Logon as a service”. If the user ID does not have the privilege to run the service, the service does not start and it returns an error in the Windows system event log. Typically, you will have run the Prepare IBM MQ wizard, and set up the user ID correctly. Only if you have configured the user ID manually is it possible that you might have a problem in Version 7.1 or later.

## Windows: IgnoredErrorCodes registry key

From Version 7.1, the registry key used to specify error codes that you do not want written to the Windows Application Event Log has changed.

The contents of this registry key are not automatically migrated. If you want to continue to ignore specific error codes, you must manually migrate the registry key.

Before Version 7.1, the key was in the following location:

```
HKLM\Software\IBM\MQSeries\CurrentVersion\IgnoredErrorCodes
```

From Version 7.1, the key is in the following location:

```
HKLM\Software\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\IgnoredErrorCodes
```

where *MQ\_INSTALLATION\_NAME* is the installation name associated with a particular installation of the product.

**Related information:**

Ignoring error codes under Windows systems

## **Windows: Installation and infrastructure information**

From Version 7.1, the location of Windows installation and infrastructure information has changed.

A top-level string value, *WorkPath*, in the *HKLM\SOFTWARE\IBM\WebSphere MQ* key, stores the location of the product data directory which is shared between all installations. The first installation on a machine specifies it, subsequent installations pick up the same location from this key.

Other information that in releases before Version 7.1 was stored in the registry on Windows is, for Version 7.1 or later, stored in *.ini* files.

**Related information:**

Changing configuration information on Windows, UNIX and Linux systems

## **Windows: Local queue performance monitoring**

On Windows, from Version 7.1, it is no longer possible to monitor local queues using the Windows performance monitor.

Instead, use the performance monitoring commands, which are common to all platforms, provided by IBM MQ.

**Related information:**

Real-time monitoring

## **Windows: “Logon as a service” required**

From Version 7.1, the user ID that runs the IBM MQ Windows service must have the user privilege to “Logon as a service”. If the user ID does not have the privilege to run the service, the service does not start and it returns an error in the Windows system event log. Typically, you will have run the Prepare IBM MQ wizard, and set up the user ID correctly. Only if you have configured the user ID manually is it possible that you might have a problem in Version 7.1 or later.

You have always been required to give the user ID that you configure to run IBM MQ the user privilege to “Logon as a service”. If you run the Prepare IBM MQ wizard, it creates a user ID with this privilege. Alternatively, it ensures that a user ID you provide has this privilege.

It is possible that you ran IBM MQ in releases before Version 7.1 with a user ID that did not have the “Logon as a service” privilege. You might have used it to configure the IBM MQ Windows service *MQSeriesServices*, without any problems. If you run an IBM MQ Windows service in Version 7.1 or later with the same user ID that does not have the “Logon as a service” privilege, the service does not start.

The IBM MQ Windows service *MQSeriesServices*, with the display name *MQSeries*, changed in Version 7.1. A single IBM MQ Windows service per server is no longer sufficient. An IBM MQ Windows service per installation is required. Each service is named *MQ\_InstallationName*, and has a display name *IBM WebSphere MQ (InstallationName)*. The change, which is necessary to run multiple installations of IBM MQ, has prevented IBM MQ running the service under a single specific user ID. In Version 7.1 or later, a *MQ\_InstallationName* service must run as a service.

The consequence is a user ID that is configured to run the Windows service *MQ\_InstallationName* must be configured to “Logon as a service”. If the user ID is not configured correctly, errors are returned in the Windows system event log.

Many installations on earlier releases, and installations from Version 7.1 onwards, configure IBM MQ with the Prepare IBM MQ wizard. The wizard sets up the user ID with the “Logon as a service” privilege

and configures the IBM MQ Windows service with this user ID. Only if, in previous releases, you have configured MQSeriesServices with another user ID that you configured manually, might you have this migration problem to fix.

### **Windows: MSCS restriction with multiple installations**

When you install or upgrade to Version 7.1 or later, the first installation of the product on the server is the only one that can be used with Microsoft Cluster Server (MSCS). No other installations on the server can be used with MSCS. This restriction limits the use of MSCS with multiple installations of the product.

When you run the **haregtyp** command it defines the first installation of the product to be installed as an MSCS resource type; see IBM MQ MSCS support utility programs. The implications are as follows:

1. You must associate queue managers that are participating in an MSCS cluster with the first installation on the server.
2. Setting the primary installation has no effect on which installation is associated with the MSCS cluster.
3. If you are upgrading from Version 7.0.1 to Version 8.0, you must follow the single-stage migration scenario; see “UNIX, Linux, and Windows: Single-stage migration from Version 7.0.1, or later, to the latest version” on page 496.

### **Windows: Migration of registry information**

Before Version 7.1, all IBM MQ configuration information, and most queue manager configuration information, was stored in the Windows registry. From Version 7.1 onwards, all queue manager configuration information (for example, `mqs.ini`, `qmstatus.ini`, and `qm.ini`) is stored in files; the same files as in UNIX and Linux platforms.

The change does not affect the operation of existing applications or queue managers, but it does affect any administrative procedures and scripts that reference the registry.

In Version 7.0.1, to support multi-instance (high availability) queue managers, the queue manager configuration information of some queue managers is stored in `qm.ini` and `qmstatus.ini` rather than in the registry.

The Version 7.0.1 configuration information is accessed from other installations. You must stop all the queue managers and IBM MQ applications running on the server to release any locks.

The transfer of configuration data from the registry to files is automatic. It takes place under either of the following circumstances:

- You migrate an existing Windows system to Version 7.1 or later.
- Version 7.0.1 is uninstalled from a server that has a Version 7.1, or later, installation.

As a consequence, after uninstallation of Version 7.0.1 on a multi-installation server, it is difficult to restore a Version 7.0.1 installation to run any queue managers that you want to restore to the `701` command level:

1. You cannot reinstall Version 7.0.1 on the server. You must run the queue managers on a different server.
2. When you transfer the queue manager data to another server, with Version 7.0.1 installed, you must create the correct registry configuration entries. The entries are not available to copy from the registry on the multi-installation server. Back up the registry entries before uninstalling Version 7.0.1.

**Related tasks:**

“UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500

Side-by-side migration is the term used to describe installing a new version of IBM MQ alongside an older version on the same server. Queue managers remain running during the installation and verification of the new version of IBM MQ. They remain associated with the older version of IBM MQ. When you decide to migrate queue managers to the new version of IBM MQ, you stop all queue managers, uninstall the old version, and migrate them all to the new version of IBM MQ.

“UNIX, Linux, and Windows: Multi-stage migration from Version 7.0.1, or later, to the latest version” on page 505

Multi-stage migration is the term used to describe running a new version of IBM MQ alongside an older version on the same server. After installing the new version alongside the old, you can create new queue managers to verify the new installation, and develop new applications. At the same time, you can migrate queue managers and their associated applications from the old version to the new. By migrating queue managers and applications one-by-one, you can reduce the peak workload on staff managing the migration.

**Related information:**

Changing configuration information on Windows, UNIX and Linux systems

**Windows: Relocation of the mqclient.ini file**

From Version 7.1, the `mqclient.ini` file has moved from `FilePath` to `WorkPath`. This is similar to the model already used on UNIX and Linux systems.

If you supply separate file and work paths, you will see a change in behavior when migrating from a release before Version 7.1. You have an additional step to perform when you choose to uninstall IBM WebSphere MQ Version 7.0 before installing Version 7.1 or later. Before uninstalling IBM WebSphere MQ Version 7.0, you must copy `mqclient.ini` directly to the `Config` directory in your data path so that it can be picked up by the Version 7.1 or later installation.

**Windows: Task manager interpretation**

From Version 7.1, the interpretation of the processes that are listed by the Windows Task Manager has changed. The results can show IBM MQ processes running for multiple installations on a server. Before Version 7.1, the process list identified IBM MQ processes running on only a single installation of the product on a Windows server.

The implications of this change depend on how the results are qualified and interpreted, and how the list of processes is used. The change affects you, only if you start to run multiple installations on a single server. If you have incorporated the list of IBM MQ processes into administrative scripts or manual procedures, you must review the usage.

On 64-bit Windows platforms, IBM MQ Version 8.0 processes and processes for earlier releases of the product can be easily identified. As queue managers for releases before IBM MQ Version 8.0 are 32 bit, all queue manager processes for releases before Version 8.0 have `'*32'` associated with their process names. In Version 8.0, the queue manager is 64 bit, hence the queue manager processes in Windows Task manager do not have an associated `'*32'`.

**Related information:**

Stopping queue managers in IBM MQ for Windows

**Windows: IBM MQ Active Directory Services Interface**

From Version 7.1, the IBM WebSphere MQ Active Directory Services Interface is no longer available.

If you are migrating from a release before Version 7.1 and your application uses the IBM WebSphere MQ Active Directory Services Interface, you must rewrite your application to use Programmable Command Formats.

**Related information:**

Introduction to Programmable Command Formats

**IBM MQ maintenance tasks (On platforms other than z/OS )**

When you apply and remove maintenance level updates to IBM MQ, no migration is required. Maintenance level updates are applied either as a fix pack, or by manually applying an interim fix.

**About this task**

This information describe how to apply and remove fix packs on platforms other than z/OS.

Follow the appropriate link in Table 100 for the platform that your enterprise uses.

*Table 100. Applying and removing maintenance*

| Apply   | Remove  |
|---------|---------|
| AIX     | AIX     |
| HP-UX   | HP-UX   |
| Linux   | Linux   |
| Solaris | Solaris |
| Windows | Windows |

**Related information:**

Getting product fixes

 [IBM WebSphere MQ](#)

**Applying and removing maintenance level updates (On platforms other than z/OS )**

When you apply and remove maintenance level updates to IBM MQ, no migration is required. Maintenance level updates are applied either as a fix pack, or by manually applying an interim fix.

**About this task**

This information describe how to apply and remove fix packs on platforms other than z/OS.

Follow the appropriate link, to either apply maintenance or restore maintenance to a previous level, for the platform, or platforms, that your enterprise uses.

## AIX: Applying maintenance level upgrades on IBM MQ

How to apply maintenance level updates to IBM MQ for AIX using `installp`.

### Before you begin

1. Ensure you have enough disk space to apply maintenance level updates. A maintenance level update requires hard disk space for installation. In addition, the installation process might require a similar amount of disk space to save the previous level. For example, a 16 MB update might require 32 MB of space. The additional space allows a maintenance level update to be removed, and the previous level to be restored automatically.
2. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

### About this task

Stop applications using the installation and use the `installp` command, to install maintenance level updates to clients and servers. Alternatively, if the installation is in the default installation location, you can use the *System Management Interface Tool*, SMIT.

**Important:** You cannot go back from a later version of the product to a prior version of the product, for example from IBM MQ Version 8.0 to IBM WebSphere MQ Version 7.x.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

Major full versions of the base product are COMMITTED by default. Fix packs on a full base version can be in APPLIED state, and you can go back one release level.

If you need the ability to revert to an earlier version, you should perform a side-by-side migration, and migrate your queue managers to the later version at any time. See “UNIX, Linux, and Windows: Side-by-side migration from Version 7.0.1, or later, to the latest version” on page 500 for further information.

However, if you start a queue manager under IBM MQ Version 8.0, that queue manager is automatically migrated, and cannot be downgraded to the previous version.

You require a backup of the IBM WebSphere MQ Version 7.x queue manager and data logs to restore if needed.

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.  
If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.
3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the `dspmqs` command to list the state of all the queue managers on the system.  
Run either of the following commands from the installation that you are updating:  

```
dspmqs -o installation -o status
dspmqs -a
```

**dspmqs -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.



**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm|sr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.

5. Install the update in one of the following ways:

- Update the whole installation in the default location:

```
installp -agXYd . all
```

- Update selected filesets in the default location:

```
installp -agXYd . list of file sets
```

- Update the whole product in a non-default location using the **-R** flag:

```
installp -R USIL_Directory -agXYd . all
```

- Update selected filesets in a non-default location using the **-R** flag:

```
installp -R USIL_Directory -agXYd . list of file sets
```

*USIL\_Directory* is the installation parent directory. IBM MQ is installed underneath the directory. For example, if `/USIL1` is specified, the IBM MQ product files are located in `/USIL1/usr/mqm`. `/USIL1/usr/mqm` is known as the *MQ\_INSTALLATION\_PATH*.

## Related information:

dspmqs

Stopping a queue manager

## AIX: Restoring the previous maintenance level on IBM MQ

How to restore a previous maintenance level using the *System Management Interface Tool* (SMIT).

### Before you begin

1. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

### About this task

You can back out maintenance updates (fix packs) and restore your system to the previous maintenance or install level, for any component of IBM MQ for AIX that is in the **APPLIED** state.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

Use the following command to display the current state of the IBM MQ for AIX filesets:

```
lsipp [-R usil] -l "mqm*"
```

To back out a maintenance update, as the user root, issue the command:

```
installp [-R usil] -r "mqm*"
```

Otherwise:

### Procedure

1. Log in as a user in group mqm.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.
3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the **dspmqs** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqs -o installation -o status
dspmqs -a
```

**dspmqs -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqs -a** displays the status of active queue managers associated with the installation from which the command is run.
  - b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```
  - c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqlsr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.
5. Open the appropriate **smiit** panel using this sequence:

```
Software Installation and Maintenance
Software Maintenance and Utilities
Reject Applied Software Updates (Use Previous Version)
```

Alternatively, use a fast path command, `smit[ty] install_update`.

6. Complete the **SOFTWARE** name field.  
Enter `mqm*` to restore all applicable file set updates to your installation.

**Note:** If an option to restore only selected file set updates for IBM MQ for AIX appears, avoid it. The option results in all applicable file set updates for the maintenance update being restored.

7. Click **Enter** to reject the current maintenance level and reinstate the previous maintenance or install level.

- a. Accept displayed default values for all other fields
- b. Dismiss the confirmation message

The reject process starts. While the command runs, it displays progress messages terminating with an **Install Summary** table.

- a. Check the table to see which components of IBM MQ for AIX have been rejected

## Related information:

dspmq

Stopping a queue manager

## HP-UX: Applying maintenance level updates on IBM MQ

How to apply maintenance level updates to IBM MQ for HP-UX using **swinstall**.

### Before you begin

1. Ensure you have enough disk space to apply maintenance level updates. A maintenance level update requires hard disk space for installation. In addition, the installation process might require a similar amount of disk space to save the previous level. For example, a 16 MB update might require 32 MB of space. The additional space allows a maintenance level update to be removed, and the previous level to be restored automatically.
2. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

### About this task

1. If you want to install both the base package and the maintenance update packages, install the base package separately first. Then install the maintenance update packages.
2. If you are using the interactive installer, click **Options> Change Options**. Then clear the **autoselect dependencies when marking software** check box before selecting the maintenance update package for installation.
3. Error messages might be seen when running **swinstall**, even when successfully updating an installation.

There are two approaches you can take to handling errors in the application of maintenance.

- a. Aim for an error-free update by applying maintenance only to those components that are installed.
- b. Apply the whole maintenance package and check the error logs, error by error, ignoring the insignificant errors.

Both approaches are described.

Many of the insignificant errors are caused by **swinstall** trying to apply updates to components that are not installed. Consider whether there are any significant errors reported with the insignificant ones.

- The following errors might not indicate a serious problem. They are written to the console, or to the **swinstall** panel.

```
ERROR: "hpux11.mycompany.com/":
The software dependencies for 15 products or filesets cannot be resolved.
```

```
ERROR: "hpux11.mycompany.com/":
17 filesets were determined to be skipped in the analysis phase.
The execution phase failed for "hpux11.mycompany.com/".
Analysis and Execution had errors.
```

- The following errors might not indicate a serious problem. They are written to the `swjob` output for a **swinstall** session.

```
ERROR: 17 of 20 filesets had Errors.
3 of 20 filesets had no Errors or Warnings.
```

```
ERROR: The Execution Phase had errors.
See the above output for details.
```

## Procedure

1. Log in as a user in group mqm.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their SYSTEM.FTE.STATE queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the **dspmqr** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqr -o installation -o status
dspmqr -a
```

**dspmqr -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, DISPLAY LSSTATUS(\*) STATUS to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm|sr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.
5. Make your current directory the location of the *Service\_update\_package*.

The file name of the *Service\_update\_package* follows the pattern `hp-Uxxxx.v11`. You must prefix *Service\_update\_package* with the absolute path to the installation file. To save typing, construct the path using the `$PWD` variable.

6. Run the HP-UX command `swlist l= MQ_INSTALLATION_PATH MQSERIES` to list all of the IBM MQ components that are installed.
7. Decide whether to install the updates interactively, and if you want to control which components are updated.

You can update in the following ways:

- Silently update all the installed IBM MQ components by installing the whole maintenance package.

```
swinstall -s $PWD/service_update_package
MQSERIES,l=MQ_INSTALLATION_PATH
```

The `swinstall` command attempts to find an installed component for every component in the update package, and updates it. `swinstall` writes out error messages for components that it cannot find.

- Silently update some IBM MQ components by installing only the required updates from the maintenance package.

If you specify `update_components` correctly, the update procedure can be error-free. `swinstall` only updates components that you have listed and components that are dependent on components you have listed.

- a. Using the list of installed IBM MQ components, create a space separated list of the components you want to update (`update_components`). This list requires the installation path of each component to be specified, in the form: `component ,l= MQ_INSTALLATION_PATH`

- b. `swinstall -s $PWD/service_update_package`  
`update_components`

- Interactively update some IBM MQ components from the maintenance package, selecting only the update components that are required.

```
swinstall -s $PWD/service_update_package
```

- a. Open **MQSERIES** and mark the update components you want to apply. Correctly marked, there are no errors when the updates are applied. The installer resolves dependencies automatically.

- b. Select **Actions > Change Product Location** to change the IBM MQ installation you intend to update.

- c. Select **Actions > Install**. The log file tells you if there are any problems that need fixing.

#### Related information:

`dspmq`

Stopping a queue manager

## HP-UX: Restoring the previous maintenance level on IBM MQ

How to restore a previous maintenance level by using `swremove`.

### Before you begin

1. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.

a. Run the **dspmqr** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqr -o installation -o status
dspmqr -a
```

**dspmqr -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

d. Stop any listeners associated with the queue managers, using the command:

```
endmqm|sr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.

5. Run the **swremove** command to remove the maintenance package from the system.

For example, to remove the 7.R.0.1 maintenance level, use the command:

```
swremove MQSERIES,r=7.R.0.1,l= MQ_INSTALLATION_PATH
```

where:

- R is the number of the Release
- `MQ_INSTALLATION_PATH` is the installation path for IBM MQ

Details of the **swremove** command can be found in the *HP-UX Administration Guide* or by using the **man swremove** command.

## Related information:

dspmq

Stopping a queue manager

## IBM i: Applying maintenance level updates on the latest release



How to apply maintenance level updates by stopping IBM MQ and using the IBM i standard maintenance procedure.

### Before you begin

To find out what version you have currently installed, use the following commands:

Table 101. IBM MQ commands to display the installed versions

| IBM MQ Product | Version command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IBM MQ Server  | DSPMQMVER                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| IBM MQ Java    | IBM MQ classes for Java:<br>java com.ibm.mq.MQJavaLevel<br><br><b>Note:</b> For this command to work, you might need to set your environment classpath to include: <ul style="list-style-type: none"><li>• /QIBM/ProdData/mqm/java/lib/com.ibm.mq.jar</li></ul> IBM MQ classes for Java Message Service:<br>java com.ibm.mq.jms.MQJMSLevel<br><br><b>Note:</b> For this command to work, you might need to set your environment classpath to include: <ul style="list-style-type: none"><li>• /QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar</li></ul> See Environment variables relevant to IBM MQ classes for Java and Environment variables relevant to IBM MQ classes for JMS. |
| IBM MQ Client  | DSPMQMVER                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

### About this task

Maintenance updates for IBM MQ for IBM i are supplied as PTFs (Program Temporary Fixes). They are available for download from the web as save files, which are normally stored in the QGPL library. IBM i PTF's can be found in "Fix Central" at the following location:

FixCentral.

### Procedure

1. Read the cover letter carefully to see if you need to take any special actions.
2. Sign on to a new interactive IBM i session, ensuring that you are not accessing any IBM MQ objects.
3. Ensure that you have:
  - a. \*ALLOBJ authority, or object management authority for the QMQM library.
  - b. Sufficient authority to use the ENDSBS command.
4. Warn all users that you are going to stop IBM MQ.
5. Use the ENDMQM command to quiesce all queue managers:



```
ENDMQM MQMNAME(*ALL) OPTION(*CNTRLD) ENDCCTJOB(*YES) RCDMQMIMG(*YES)
TIMEOUT(15)
```

Where 15 is a timeout value in seconds.

If the ENDMQM command has not completed within a reasonable period (at least 10 minutes), use the WRKMQM command. This command identifies the queue managers that are still ending. Then force each one in turn to stop by issuing:

```
ENDMQM MQMNAME(QMGRNAME) OPTION(*IMMED)
```

Where *QMGRNAME* is the name of the queue manager.

Complete the tidying up of shared memory by issuing the command:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED) ENDCCTJOB(*YES) RCDMQMIMG(*NO)
TIMEOUT(15)
```

6. If the command in the previous step does not complete, end the subsystem immediately by issuing:  
ENDSBS SBS(QMQM) OPTION(\*IMMED)
7. If this command also fails, use the operating system command ENDJOB to end all jobs in the subsystem QMQM, as described in the following steps.

**Note:** Do not use ENDJOBABN unless you intend to perform an IPL on the machine before starting IBM MQ. Ending IBM MQ jobs using ENDJOBABN can lead to damaged semaphores, which in turn can prevent your queue manager from starting.

- a. If a QMGR must be shut down manually, the recommended order of ending jobs (ENDJOB) is shown in the list that follows. Wait a few minutes for AMQA\* or AMQZ\* jobs to tidy up.

- 1) RUNMQLSR - TCP listener (multi-threaded)
- 2) AMQCLMAA - TCP listener (single-threaded)
- 3) AMQRMPPA - Channel process pooling job
- 4) RUNMQCHI - channel initiator
- 5) AMQCRSTA - receiving MCA jobs
- 6) RUNMQCHL - sending MCA jobs
- 7) AMQCRS6B - LU62 receiver channel
- 8) AMQPCSEA - command server
- 9) RUNMQTRM - Application trigger monitor
- 10) RUNMQDLQ - Dead letter queue handler
- 11) AMQFCXBA - IBM Integration Bus Worker Job
- 12) AMQFQPUB - Queued Publish/Subscribe Daemon
- 13) RUNMQBRK - IBM Integration Bus Control Job
- 14) AMQZMUC0 ('0' is a zero) - Utility Manager
- 15) AMQZMUF0 ('0' is a zero) - Utility Manager
- 16) AMQZMUR0 ('0' is a zero) - Utility Manager
- 17) AMQZMGR0 ('0' is a zero) - Process Controller
- 18) AMQRRMFA - cluster repository manager
- 19) AMQZDMAA - deferred message manager
- 20) AMQALMPX - Log Manager
- 21) AMQZFUMA - object authority manager
- 22) AMQZLSA0 ('0' is a zero) - LQM agents
- 23) AMQZLAA0 ('0' is a zero) - LQM agents
- 24) AMQZXMA0 ('0' is a zero) - Execution Controller

- b. Issue the following command:

```
ENDMQM MQMNAME(QMGRNAME) OPTION(*IMMED)
```

- c. Issue the following command:

```
ENDMQM MQMNAME(*ALL) OPTION(*CNTRLD) ENDCCTJOB(*YES) RCDMQMIMG(*NO)
TIMEOUT(05)
```

Where 05 is a timeout value in seconds.

- d. Manually clean up shared memory. Issue the following command:

```
EDTF '/QIBM/UserData/mqm/qmgrs'
```

then:

- 1) Take option 5 for **&SYSTEM** and check that the following directories are empty: `isem`, `esem`, `msem`, `ssem`, and `shmem`.
- 2) Take option 5 for **QMGRNAME** and check that the following directories are empty:- `isem`, `esem`, `msem`, `ssem`, and `shmem`.
- 3) Take option 5 for **&ipcc** in the QMGRNAME directory and check that the following directories are empty:- `isem`, `esem`, `msem`, `ssem`, and `shmem`.
- 4) Take option 5 for **&qmpersist** in the QMGRNAME directory and check that the following directories are empty:- `isem`, `esem`, `msem`, `ssem`, and `shmem`.
- 5) Take option 5 for **&app** and check that the following directories are empty: `isem`, `esem`, `msem`, `ssem`, and `shmem`.

8. Load and apply a PTF

## IBM i: Restoring a queue manager from the latest release to a previous release

You can restore a queue manager to the previous version of the product from the latest version, if you have made a backup of the system or queue manager. If you have started the queue manager and processed any messages, or changed the configuration, the task cannot give you any guidance on restoring the current state of the queue manager.

### Before you begin

1. You must have made a backup of the system or queue manager before you upgraded to the latest version. For more information see [Backing up and restoring IBM MQ queue manager data](#)
2. If any messages were processed after starting the queue manager, you cannot easily undo the effects of processing the messages. You cannot restore the queue manager to the previous version of the product in its current state. The task cannot give you any guidance how to deal with subsequent changes that have occurred. For example, messages that were indoubt in a channel, or in a transmission queue on another queue manager, might have been processed. If the queue manager is part of a cluster, then configuration messages and application messages might have been exchanged.

### About this task

When you restore a previous version of a queue manager, you restore the queue manager to its earlier code level. Queue manager data is restored to the state it was in when the queue manager was backed up.

### Procedure

1. Stop the queue manager.
2. If you performed a slip install, you must reinstall IBM MQ.
  - a. Uninstall the previous installation.
  - b. Reinstall IBM MQ from a manufacturing refresh.
  - c. Apply the fix pack and interim fixes that restore IBM MQ to its previous level.
  - d. Restore the queue manager data from the backup taken before installing the latest version.
3. Restart the previous version queue manager.

## Related information:

Backing up and restoring a queue manager

## Linux: Applying maintenance level updates on IBM MQ

How to apply maintenance level updates to IBM MQ for Linux using RPM. The following procedure applies to all Linux platforms, including Ubuntu.

### Before you begin

If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

### About this task

Maintenance level updates are delivered in the form of Red Hat Package Manager (RPM) update images, which are applied using the RPM installation tool.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

**Important:** `pax` and `rpmbuild` are not supplied as part of the product. You must obtain these from your Linux distribution supplier.

Additional disk space is required for the update images to allow maintenance level updates to be removed and the previous level restored. The updated files are kept in `MQ_INSTALLATION_PATH/maintenance` directory. Do not delete or move this directory or the files it contains.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Updates are cumulative. You can apply your chosen update directly, without applying any previous updates first. The maintenance level updates might contain updates for one or more packages. You must apply those parts of an update that correspond to the packages that are applied in your installation.

**Important:** Although it is possible to install a fix pack at the same level as an installation performed from a manufacturing refresh image at that level, you should not attempt this process. Installing a fix pack at the same level as the one already on your system, can leave the package management database of your system in an inconsistent state with respect to the installation of IBM MQ.

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the `dspmqs` command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqs -o installation -o status
dspmqs -a
```

**dspmqs -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm lsr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.
5. Change into the directory containing the maintenance packages.
6. Run the **ls** command to list the available updates.

For example, if there are level 1 maintenance updates for the Runtime, SDK and Server packages, you see the following:

```
MQSeriesRuntime-Uxxx-V.R.0-1.i386.rpm
MQSeriesSDK-Uxxx-V.R.0-1.i386.rpm
MQSeriesServer-Uxxx-V.R.0-1.i386.rpm
```

where *V* is the version number and *R* is the number of the Release.

7. Run the **rpm** command to find out which packages are installed on your server.

Enter the following command:

```
rpm -qa | grep MQSeries
```

#### Notes:

- a. If you are using Ubuntu, add the **--force-debian** attribute.

```
rpm --force-debian -qa | grep MQSeries
```

- b. If you are using Linux on POWER Systems - Little Endian, add the **--ignorearch** attribute.

You must include this option to prevent problems with some levels of rpm not recognizing the Linux on POWER Systems - Little Endian architecture.

For example, if you have a minimum IBM MQ installation and SDK component, at level 0, the **rpm** command returns:

```
MQSeriesRuntime-V.R.0-0
MQSeriesSDK-V.R.0-0
MQSeriesServer-V.R.0-0
```

where V is the version number and R is the number of the Release.

8. If this fix pack is to be upgraded on an installation, other than the first installation on the system, run the **crtmqfp** command to create and use a unique set of packages to install on the system. Note, that if this is the first, or only, IBM MQ installation on the system, you can ignore this step. You must install the **pax** command in order for the **crtmqfp** command to run on Linux.

- a. Run the command `./crtmqfp <suffixname>` where *suffixname* is the same as the suffix used during renaming of the base level IBM MQ installation.

- b. Set your current directory to the location specified when the **crtmqfp** command completes. This directory is a subdirectory of `/var/tmp/mq_rpms`, in which the unique set of packages is created. The packages have the suffix value contained within the filename.

For example, if you used suffix 1 during repackaging of the base level IBM MQ installation, enter the command: `./crtmqfp 1`.

There is now a subdirectory named `/var/tmp/mq_rpms/1/xxxx`, and the packages will be renamed, for example, from `MQSeriesRuntime-V.R.0-1.xxxx.rpm` to `MQSeriesRuntime_1-V.R.0-1.xxxx.rpm`. Where V is the version number and R is the number of the Release.

9. Run the **rpm** command to apply all available updates for the packages you have on your system:

- To update an installation in the default location, `/opt/mqm`:

```
rpm -ivh MQSeriesRuntime-Uxxxx-V.R.0-1.i386.rpm
 MQSeriesSDK-Uxxxx-V.R.0-1.i386.rpm
 MQSeriesServer-Uxxxx-V.R.0-1.i386.rpm
```

where V is the version number and R is the number of the Release.

- To update an installation in a custom location, specify the **rpm** prefix option:

```
rpm --prefix /opt/customLocation -ivh MQSeriesRuntime-Uxxxx-V.R.0-1.i386.rpm
 MQSeriesSDK-Uxxxx-V.R.0-1.i386.rpm
 MQSeriesServer-Uxxxx-V.R.0-1.i386.rpm
```

where V is the version number and R is the number of the Release.

You must apply all packages in a maintenance update that correspond to those packages that are currently installed on your system.

10. Repeat step 7 on page 636 to list the packages that are now available.

The Runtime, SDK, and Server packages are now at level 1:

```
MQSeriesRuntime-V.R.0-0
MQSeriesSDK-V.R.0-0
MQSeriesServer-V.R.0-0
MQSeriesRuntime-Uxxxx-V.R.0-1
MQSeriesSDK-Uxxxx-V.R.0-1
MQSeriesServer-Uxxxx-V.R.0-1
```

where V is the version number and R is the number of the Release.

**Note:**

After the installation of IBM MQ fix packs, if you run the `rpm-verify` or `rpm -V` command, it does not return the correct results. It produces spurious results relating to missing files in `MQ_INSTALLATION_PATH/maintenance`.

This error message can be ignored because it is a known limitation in the IBM MQ fix pack installation code. For further information about this error, see IBM MQ Fix Pack install errors - Linux reports errors

## What to do next

For further information about using RPM to install software packages, see your Linux documentation.

### Related information:

`dspmqr`

Stopping a queue manager

## Linux: Restoring the previous maintenance level on IBM MQ

How to remove updates and restore the previous maintenance level using **RPM**. The following instructions apply to all Linux platforms, including Ubuntu.

### Before you begin

If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

### About this task

When maintenance is applied, the original versions of replaced files are saved to allow the updates to be removed if necessary. To restore the previous maintenance level, run an Red Hat Package Manager, RPM, `uninstall` command for all the packages that were updated by the maintenance package as follows:

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.
3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the **`dspmqr`** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqr -o installation -o status
dspmqr -a
```

**`dspmqr -o installation -o status`** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**`dspmqr -a`** displays the status of active queue managers associated with the installation from which the command is run.
  - b. Run the **`MQSC`** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```
  - c. Run the **`endmqm`** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm lsr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.
5. Run the **rpm** command to find out which packages are installed on your server.

Enter the following command:

```
rpm -qa | grep MQSeries
```

**Note:** If you are using Ubuntu, add the **--force-debian** attribute.

```
rpm --force-debian -qa | grep MQSeries
```

Using the example given in “Linux: Applying maintenance level updates on IBM MQ” on page 635, returns:

```
MQSeriesRuntime-V.R.0-0
MQSeriesSDK-V.R.0-0
MQSeriesServer-V.R.0-0
MQSeriesRuntime-Uxxx-V.R.0-1
MQSeriesSDK-Uxxx-V.R.0-1
MQSeriesServer-Uxxx-V.R.0-1
```

where V is the version number and R is the number of the Release.

6. Run the **rpm** command to remove all the updates applied at level 1.

Enter the following commands:

```
rpm -ev MQSeriesRuntime-Uxxx-V.R.0-1 MQSeriesSDK-Uxxx-V.R.0-1
MQSeriesServer-Uxxx-V.R.0-1
```

where V is the version number and R is the number of the Release.

7. Repeat step 5 to check that the ptf packages have been removed, leaving only the original install packages:

```
MQSeriesRuntime-V.R.0-0
MQSeriesSDK-V.R.0-0
MQSeriesServer-V.R.0-0
```

where V is the version number and R is the number of the Release.

## What to do next

For further information on using RPM to install software packages, see your Linux documentation.

### Related information:

`dspmqr`

Stopping a queue manager

## Solaris: Applying maintenance level updates on IBM MQ

How to apply maintenance level updates to IBM MQ for Solaris using `pkgadd`.

### Before you begin

1. Ensure you have enough disk space to apply maintenance level updates. A maintenance level update requires hard disk space for installation. In addition, the installation process might require a similar amount of disk space to save the previous level. For example, a 16 MB update might require 32 MB of space. The additional space allows a maintenance level update to be removed, and the previous level to be restored automatically.
2. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

### About this task

Stop applications using the installation and use `pkgadd` to install maintenance.

**Important:** Although it is possible to install a fix pack at the same level as an installation performed from a manufacturing refresh image at that level, you should not attempt this process. Installing a fix pack at the same level as the one already on your system, can leave the package management database of your system in an inconsistent state with respect to the installation of IBM MQ.

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.  
If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.
3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the `dspmqr` command to list the state of all the queue managers on the system.  
Run either of the following commands from the installation that you are updating:  

```
dspmqr -o installation -o status
dspmqr -a
```


**dspmqr -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.
  - b. Run the `MQSC` command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.



```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm lsr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.
5. Change into the directory containing the maintenance packages.
6. Run the **crtmqfp** command to create and use a unique set of packages to install on the system, if this fix pack is to be upgraded on a installation that is not the first installation on the system. This command creates and uses a unique set of packages to install on the system.
  - a. Run the command **crtmqfp** **mqm-** <suffixname> where *suffixname* is the same as the suffix used during renaming of the base level IBM MQ installation. Note that this command creates a full copy of the installation packages in a subdirectory of **/var/tmp**.
  - b. Set your current directory to the location specified when the **crtmqfp** command completes. This directory is a subdirectory of **/var/spool**, in which the unique set of packages is created. The packages have the suffix value contained within the filename.
7. Proceed with installation using the following command: Enter the following command to start the installation process if this fix pack is to be upgraded on an installation that is
  - a. The first installation on the system:

```
pkgadd -d packagename
```

where *packagename* corresponds to the image file name. For example:

```
mqm-U1234.img
```

- b. Not the first installation on the system:

```
pkgadd mqm-suffixname
```

where *suffixname* is the name of the directory created in **/var/spool/pkg**.

For example, if you install IBM WebSphere MQ Version 7.0 as a package called `mqm-main7` and create a package to upgrade to IBM WebSphere MQ Version 7.0.0.1, using the command `crtmqfp mqm-main7`, package `mqm-main7-07-00-00-01` is created in `/var/spool/pkg`.

To install package `mqm-main7-07-00-00-01`, issue the command `pkgadd mqm-main7-07-00-00-01`.

For further information about using `pkgadd` to install software packages, see the Solaris documentation.

8. Follow the on-screen instructions.

#### Related information:

`dspmq`

Stopping a queue manager

## Solaris: Applying maintenance level updates on IBM MQ in non-interactive mode

You can install IBM MQ for Solaris non-interactively by creating a response file and an admin file.

### Before you begin

1. Ensure you have enough disk space to apply maintenance level updates. A maintenance level update requires hard disk space for installation. In addition, the installation process might require a similar amount of disk space to save the previous level. For example, a 16 MB update might require 32 MB of space. The additional space allows a maintenance level update to be removed, and the previous level to be restored automatically.
2. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

You can apply and remove maintenance from a IBM MQ MQI client that is not installed on the same server as a queue manager. You do not have to stop any queue managers or logon as administrator. Because you do not have to stop any queue managers, do not do steps 1 to 3 in the following maintenance procedure.

### About this task

Stop applications using the installation and use `pkgadd` to install maintenance.

**Important:** Although it is possible to install a fix pack at the same level as an installation performed from a manufacturing refresh image at that level, you should not attempt this process. Installing a fix pack at the same level as the one already on your system, can leave the package management database of your system in an inconsistent state with respect to the installation of IBM MQ.

### Procedure

1. Log in as a user in group `mqm`.
2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the `dspmq` command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmq -o installation -o status
dspmq -a
```

`dspmq -o installation -o status` displays the installation name and status of queue managers associated with all installations of IBM MQ.

`dspmq -a` displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.
 

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```
- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:
 

```
endmqm lsr -m QMgrName
```
4. Log in as root, or switch to the superuser using the **su** command.
5. Change into the directory containing the maintenance packages.
6. Run the **crtmqfp** command to create and use a unique set of packages to install on the system, if this fix pack is to be upgraded on a installation that is not the first installation on the system. This command creates and uses a unique set of packages to install on the system.
  - a. Run the command **crtmqfp** `mqm- <suffixname>` where *suffixname* is the same as the suffix used during renaming of the base level IBM MQ installation. Note that this command creates a full copy of the installation packages in a subdirectory of `/var/tmp`.
  - b. Set your current directory to the location specified when the **crtmqfp** command completes. This directory is a subdirectory of `/var/spool`, in which the unique set of packages is created. The packages have the suffix value contained within the filename.
7. Create the non-interactive install response file using the **pkgask** command. Enter the following command to create the response file if this fix pack is to be upgraded on an installation that is:
  - a. The first installation on the system:
 

```
pkgask -d <location_to_image>/imagefile -r response.txt packagename
```

 where *imagefile* corresponds to the image file name, for example `mqm-U200403.img`, *response.txt* is the name of the response file to create, and *packagename* is the fix pack package name, for example `mqm-07-05-00-02`.

- b. Not the first installation on the system:

```
pkgask -d /var/spool/pkg -r response.txt mqm-suffixname
```

where `/var/spool/pkg` is the location of the new package, `response.txt` is the name of the response file to create, and `suffixname` is the name of the directory created in `/var/spool/pkg`.

- Find the `admin_file` from the server installation media located at `<install_media>/silent/admin` or create an `admin_file` in the following format:

```
mail=
instance=unique
partial=ask
runlevel=ask
idepend=ask
rdepend=ask
space=ask
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

- Run the `pkgadd` command to apply the maintenance level update IBM MQ for Solaris in non-interactive mode. Enter the following command to start the installation process if this fix pack is to be upgraded on an installation that is:

- The first installation on the system:

```
pkgadd -v -n -r response.txt -a admin_file -d <location_to_image>/imagefile packagename
```

where `admin_file` is a path qualified name of the admin file you created, and `packagename` corresponds to the fix pack package being installed.

- Not the first installation on the system:

```
pkgadd -v -n -r response.txt -a admin_file -d /var/spool/pkg mqm-suffixname
```

- Follow the on-screen instructions.

#### Related information:

`dspmqr`

Stopping a queue manager

## Solaris: Restoring the previous maintenance level on IBM MQ

How to restore a previous maintenance level by stopping IBM MQ and using `pkggrm`.

### Before you begin

If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.

### About this task

When maintenance is applied, the original versions of replaced files are saved to allow the updates to be removed if necessary. To restore the previous maintenance level, run `pkggrm` command for all the packages that were updated by the maintenance package as follows:

### Procedure

- Log in as a user in group `mqm`.

- Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.

- End all the activity of queue managers associated with the IBM MQ installation.

- Run the `dspmqr` command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmqr -o installation -o status
dspmqr -a
```

**dspmqr -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmqr -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo "DISPLAY LSSTATUS(*) STATUS" | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:

```
endmqm|sr -m QMgrName
```

4. Log in as root, or switch to the superuser using the **su** command.

5. Run the **pkgrm** command to remove the latest maintenance update from the system:

```
pkgrm packagename
```

*packagename* is the name of the package that you want to remove; for example, `mqm-07-R-00-01`, where R is the number of the Release.

Details of the **pkgrm** command can be found in the Solaris documentation, or by using the **man pkgrm** command.

If you do not know the name of the package to remove, try listing the packages that are installed using the following command: `pkginfo | grep mqm`

**Note:** Ignore any error messages of the form `<shared pathname not removed>`.

## What to do next

If you have installed an IBM MQ MQI client, and the client was updated after installing the maintenance level that is being removed, you must specifically update your IBM MQ MQI client installation again,

after the maintenance level has been removed

**Related information:**

dspmq

Stopping a queue manager

## Windows: Applying maintenance level upgrades on IBM MQ servers

How to apply maintenance level updates to IBM MQ for Windows.

### Before you begin

1. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.
2. Download the maintenance package from the IBM MQ Support website.
3. If User Account Control (UAC) is enabled, the user who does the installation must have Administrative authority. You must elevate any command or command prompt by selecting **Run as Administrator**. If you do not, the error AMQ4353 is written in the installation log

### Procedure

1. Log on as an Administrator.

2. Stop all applications using the IBM MQ installation.

If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.

3. End all the activity of queue managers associated with the IBM MQ installation.

- a. Run the **dspmq** command to list the state of all the queue managers on the system.

Run either of the following commands from the installation that you are updating:

```
dspmq -o installation -o status
dspmq -a
```

**dspmq -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.

**dspmq -a** displays the status of active queue managers associated with the installation from which the command is run.

- b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.

```
echo DISPLAY LSSTATUS(*) STATUS | runmqsc QmgrName
```

- c. Run the **endmqm** command to stop each running queue manager associated with this installation.



The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:
 

```
endmq1sr -m QMgrName
```
4. Stop the IBM MQ service for the installation.
  - a. Right-click the **WebSphere MQ** icon in the taskbar > click **Stop WebSphere MQ**.
5. Load and apply the maintenance files for server installations:
  - Interactively:
    - a. Open the folder where the maintenance package has been extracted.
    - b. Right-click on the maintenance program and select **Run as administrator** to start the loading process.
    - c. Choose your installation language, and click **OK**.
    - d. Continue to follow the instructions on screen.
 

If you choose to load the files without applying them to an installation, you can apply the files later, as described in step 6 on page 648
  - Silently:
    - a. Open the folder where the maintenance package has been extracted.
    - b. Modify the response file, `silent_install.resp`. For details on the properties you can specify in the response file, see Table 102

Table 102. Properties used to install or uninstall a maintenance update

| Property            | Value                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQPLOG              | <i>path\file_name</i>    | <p>Pass a valid path to specify the log to be used during installation/uninstallation, for example MQPLOG="C:\TEMP\UPDATEINSTALL.LOG"</p> <p>If MQPLOG is not specified (which is the case if you start maintenance by clicking on the <b>Apply fix pack n.n.n.n</b> icon in the IBM WebSphere MQ program group) the log name used by default will be <code>amqicsdn.txt</code> in your TEMP directory ( %TEMP%).</p>                                                                                                                                                            |
| MQPINSTALLATIONNAME | <i>Installation name</i> | <p>The name of the installation to maintain. If there is only 1 installation (of any level) on the machine, this argument can be safely omitted.</p> <p>If there is more than 1 installation on the machine, <code>amqicsdn.exe</code> checks the value of MQPINSTALLATIONNAME. If one is not supplied, or the one that is supplied is unsuitable, then a GUI selection box appears. This selection box provides a list of installations to which this fix pack is applicable. If none are applicable, then <code>amqicsdn.exe</code> issues error message AMQ4781 and ends.</p> |
| MQPBACKUPPATH       | <i>path</i>              | <p>Specifies the directory to back up into during installation, for example MQPBACKUPPATH="C:\BACKUP"</p> <p>The directory, and any intermediate directories, you specify must already exist. If any one of the directories does not already exist, the install fails.</p>                                                                                                                                                                                                                                                                                                       |

Table 102. Properties used to install or uninstall a maintenance update (continued)

| Property   | Value | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQPREBOOT  | 0 1   | Specifies what to do when a reboot is required, for example MQPREBOOT=1.<br><br>If no value is supplied, you are prompted for what to do next.<br><br>If MQPREBOOT is set to 0, a reboot is suppressed<br><br>If MQPREBOOT is set to 1, the reboots go ahead without prompting.                                                                                                                                                                                      |
| MQPINUSEOK | 0 1   | Specifies whether to continue even if a file is found to be currently locked by another application. If you choose to continue even if a file is locked by another application, then you must reboot to complete fix pack installation.<br><br>If no value is supplied, or if MQPINUSEOK is set to 0, the installation fails if files are found to be in use by other applications.<br><br>If MQPINUSEOK is set to 1, the installation is deferred until you reboot. |

c. Open an elevated command prompt in the directory where the maintenance program was extracted.

d. Start the silent loading by entering the following command:

```
executableName -f responseFile
```

where:

- *executableName* is the name of the maintenance package. For example, for Version 7.1, fix pack 1: 7.1.0-WS-MQ-Windows-FP0001.exe.
- *responseFile* is the full path and name of the response file.

6. Optional: Apply the maintenance to other server installations on the system:

- Interactively:

a. From the Windows start menu, select **Start > Programs > IBM WebSphere MQ > Apply Fix Pack <V.R.M.L>**.

where

- V is the version number
- R is the release number
- M is the modification number
- L is the level of modification

b. Continue to follow the instructions on screen.

- Silently:

a. Open an elevated command prompt and navigate to the directory where the maintenance program was loaded. By default, the path is C:\Program Files (x86)\IBM\source\WebSphere MQ <V.R.M.L>

where

- V is the version number
- R is the release number
- M is the modification number
- L is the level of modification

b. Enter the following command:

```
amqicsdn MQPINSTALLATIONNAME= name MQPSILENT=1
```



where *name* is the name of the installation that you want to apply maintenance to.  
You can add other properties to the command, as listed in Table 102 on page 647.

## What to do next

On a server installation, you must restart the IBM WebSphere MQ taskbar application manually after the maintenance application completes.

If the IBM WebSphere MQ service is manually stopped, it must be manually restarted on the server. If the IBM WebSphere MQ service is not manually stopped, it is restarted automatically on the server.

The taskbar application is not restarted for any logged in sessions. Start the taskbar application in one of three ways:

1. Start the taskbar application manually from the start menu.
2. Log off and log back on again.
3. Run the command:

```
MQ_INSTALLATION_PATH\bin64\amqmtbrn.exe -Startup
```

### Related information:

dspmqr

Stopping a queue manager

## Windows: Applying maintenance level upgrades on IBM MQ clients

How to apply maintenance level updates to IBM MQ for Windows clients.

### About this task

You can apply and remove maintenance from an IBM MQ client interactively or by using the **msiexec** command to perform a silent MSI upgrade.

#### Interactive client upgrade

On the client installation media, navigate to the \Windows\MSI\ directory, then run the Setup.exe file.

#### Silent client upgrade

As an alternative method for applying maintenance to IBM MQ client systems, you can use the command **msiexec** from the command line to perform a silent MSI upgrade.

To upgrade a computer with only a single installation, you can use a command similar to the following example:

```
msiexec /i "PATH\Windows\MSI\IBM WebSphere MQ.msi" /l*v <install_log_path>
/q TRANSFORMS="1033.mst" REINSTALL=ALL REINSTALLMODE=vomus
```

For a multi installation computer with multiple clients, you can upgrade a single client by using a command similar to the following example:

```
msiexec /i "PATH\Windows\MSI\IBM WebSphere MQ.msi" /l*v <install_log_path>
/q TRANSFORMS=":InstanceId2.mst;1033.mst" REINSTALL=ALL REINSTALLMODE=vomus
```

## Results

When the maintenance completes you can query the maintenance level by running the **dspmqr** command. For more details, see “Querying the maintenance level” on page 669.

## Windows: Restoring the previous backup version on IBM MQ

How to remove updates and restore the previous maintenance level using the Windows installer

### Before you begin

1. If you are running on a server with multiple IBM MQ installations, you must identify the installation. Make sure that the commands you enter run against the correct installation; see `setmqenv`.
2. If User Account Control (UAC) is enabled, the user who does the installation must have Administrative authority. You must elevate any command or command prompt by selecting **Run as Administrator**. If you do not, the error AMQ4353 is written in the installation log

### About this task

If you applied maintenance to IBM MQ, you can restore IBM MQ to a previous level of maintenance. If you installed IBM MQ at a particular maintenance level, a *Manufacturing Refresh*, you cannot restore IBM MQ to an earlier maintenance level.

### Procedure

1. Log on as an Administrator.
2. Stop all applications using the IBM MQ installation.  
If you use the MQ Managed File Transfer (MFT) component, ensure that any MFT agents have finished all of the file transfers that they were engaged in. There should be no incomplete transfers associated with the agents, and their `SYSTEM.FTE.STATE` queues should contain no messages.
3. End all the activity of queue managers associated with the IBM MQ installation.
  - a. Run the **dspmq** command to list the state of all the queue managers on the system.  
Run either of the following commands from the installation that you are updating:  

```
dspmq -o installation -o status
dspmq -a
```

**dspmq -o installation -o status** displays the installation name and status of queue managers associated with all installations of IBM MQ.  
**dspmq -a** displays the status of active queue managers associated with the installation from which the command is run.
  - b. Run the **MQSC** command, `DISPLAY LSSTATUS(*) STATUS` to list the status of listeners associated with a queue manager.  

```
echo DISPLAY LSSTATUS(*) STATUS | runmqsc QmgrName
```
  - c. Run the **endmqm** command to stop each running queue manager associated with this installation.

```
endmqm [-c] [-w] [-i] [-p] QmgrName
```

The **endmqm** command informs an application that the queue manager it is connected to is stopping; see [Stopping a queue manager](#).

For the maintenance to proceed, applications must respond to an **endmqm** command by disconnecting from the queue manager and releasing any IBM MQ libraries they have loaded. If they do not, you must find another way to force applications to release IBM MQ resources, such as by stopping the applications.

You must also stop applications that are using the client libraries that are part of the installation. Client applications might be connected to a different queue manager, running a different installation of IBM MQ. The application is not informed about queue managers in the current installation being shut down.

Any applications that continue to have IBM MQ shared libraries from the installation loaded prevent you applying IBM MQ maintenance. An application might disconnect from a queue manager, or be forcibly disconnected, but keep an IBM MQ shared library loaded.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- d. Stop any listeners associated with the queue managers, using the command:
 

```
endmq1sr -m QMgrName
```
4. Stop the IBM MQ service for the installation.
  - a. Right-click the **WebSphere MQ** icon in the taskbar > click **Stop WebSphere MQ**.
5. Remove the maintenance interactively, or silently using a command.
  - Interactively:
    - a. For each installation of IBM MQ that has had maintenance applied, you are presented with one of the following icons in the Windows start menu:
      - 1) **Start > Programs > IBM WebSphere MQ > Remove Refresh Pack <V.R.M.L> (installation name)**
      - 2) **Start > Programs > IBM WebSphere MQ > Remove Fix Pack <V.R.M.L> (installation name)**
 where
      - V is the version number
      - R is the release number
      - M is the modification number
      - L is the level of modification
    - b. Select the installation you want to maintain and click **Remove** to start the process. This returns the installation to the state it was in before the maintenance package was applied.
  - Silently:
    - a. Open an elevated command prompt and enter the following command:
 

```
amqicsdn.exe MQPINSTALLATIONNAME= name MQPUNINST=1 MQPSILENT=1
```

 where *name* is the name of the installation that you want to remove maintenance from. You can add other properties to the command, as listed in Table 103.

Table 103. Properties used to install or uninstall a maintenance update

| Property | Value                 | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|----------|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQPLOG   | <i>path\file_name</i> | <p>Pass a valid path to specify the log to be used during installation/uninstallation, for example MQPLOG="C:\TEMP\UPDATEINSTALL.LOG"</p> <p>If MQPLOG is not specified (which is the case if you start maintenance by clicking on the <b>Apply fix pack n.n.n.n</b> icon in the IBM WebSphere MQ program group) the log name used by default will be amqicsdn.txt in your TEMP directory ( %TEMP% ).</p> |

Table 103. Properties used to install or uninstall a maintenance update (continued)

| Property            | Value                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQPINSTALLATIONNAME | <i>Installation name</i> | The name of the installation to maintain. If there is only 1 installation (of any level) on the machine, this argument can be safely omitted.<br><br>If there is more than 1 installation on the machine, amqicsdn.exe checks the value of MQPINSTALLATIONNAME. If one is not supplied, or the one that is supplied is unsuitable, then a GUI selection box appears. This selection box provides a list of installations to which this fix pack is applicable. If none are applicable, then amqicsdn.exe issues error message AMQ4781 and ends. |
| MQPBACKUPPATH       | <i>path</i>              | Specifies the directory to back up into during installation, for example MQPBACKUPPATH="C:\BACKUP"<br><br>The directory, and any intermediate directories, you specify must already exist. If any one of the directories does not already exist, the install fails.                                                                                                                                                                                                                                                                             |
| MQPREBOOT           | 0 1                      | Specifies what to do when a reboot is required, for example MQPREBOOT=1.<br><br>If no value is supplied, you are prompted for what to do next.<br>If MQPREBOOT is set to 0, a reboot is suppressed<br>If MQPREBOOT is set to 1, the reboots go ahead without prompting.                                                                                                                                                                                                                                                                         |
| MQPINUSEOK          | 0 1                      | Specifies whether to continue even if a file is found to be currently locked by another application. If you choose to continue even if a file is locked by another application, then you must reboot to complete fix pack installation.<br><br>If no value is supplied, or if MQPINUSEOK is set to 0, the installation fails if files are found to be in use by other applications.<br>If MQPINUSEOK is set to 1, the installation is deferred until you reboot.                                                                                |

6. Optional: If you no longer need the maintenance files that were loaded onto the system before maintenance was applied, you can remove them using **Add/Remove programs** or **Programs and Features** from the Control Panel. If you want to remove a maintenance file silently, run the following command:

```
<patch_install_files>_WebSphere MQ (fix pack <V.R.M.L> files)_installation\Change WebSphere MQ (fix pack <V.R.M.L> files) Installation.exe" -i silent
```

where <patch\_install\_files> is the installation directory where maintenance files are installed.

By default, this directory is C:\Program Files (x86)\IBM\source\WebSphere MQ <V.R.M.L>

**Notes:**

- a. Run the command from outside the directory, otherwise the directory is not removed.
- b. If you omit **-i silent**, the command initiates the Graphical User Interface installer.

**What to do next**

On a server installation, you must restart the IBM WebSphere MQ taskbar application manually after the maintenance application completes.

If the IBM WebSphere MQ service is manually stopped, it must be manually restarted on the server. If the IBM WebSphere MQ service is not manually stopped, it is restarted automatically on the server.

The taskbar application is not restarted for any logged in sessions. Start the taskbar application in one of three ways:

1. Start the taskbar application manually from the start menu.
2. Log off and log back on again.
3. Run the command:

```
MQ_INSTALLATION_PATH\bin64\amqmtbrn.exe -Startup
```

#### Related information:

dspmq

Stopping a queue manager

## Applying maintenance level upgrades to multi-instance queue managers

Use multi-instance queue managers to reduce the outage caused by applying maintenance updates. Follow these steps to apply maintenance to a multi-instance queue manager.

### Before you begin

Maintenance is applied to the IBM MQ installation on a server and not to individual queue managers. You must stop all the queue managers, and any IBM MQ service, on a server before you apply maintenance.

If you want a queue manager to keep running while maintenance is applied, you must configure it as a *multi-instance* queue manager, and have a standby instance running on another server. If a queue manager is an existing single instance queue manager, you must convert it to a multi-instance queue manager. See the topic, Multi-instance queue managers for prerequisites and guidance how to create a multi-instance queue manager.

You can create a multi-instance queue manager from v7.0.1 onwards. If you are running multi-instance queue managers, you then can apply a maintenance update to a *running* queue manager by switching the active instance to a different server.

Typically active and standby installations are maintained at the same maintenance level. Consult the maintenance instructions for each upgrade. Consult the instructions to see if it is possible to run the active and standby instances at different maintenance levels. Check whether fail over from higher to lower, or only lower to higher maintenance level is possible.

The instructions for applying a maintenance update might require you to stop a multi-instance queue manager completely.

If you have a primary server for running active queue manager instances, and a secondary server that runs standby instances, you have a choice of updating the primary or secondary server first. If you update the secondary server first, you must switch back to the primary server when both servers have been updated.

If you have active and standby instances on several servers, you must plan in what order you update the servers to minimize the disruption caused by ending the active instances on each server you update.

## About this task


Combine the steps in this task with the maintenance update procedure for applying maintenance to an IBM MQ server installation.

### Procedure


1. Where the maintenance update procedure instructs you to stop all running queue managers, or quiesce IBM MQ do the following instead:

The maintenance update procedure varies by platform; see “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.

- a. If the queue manager is running as standby:

-  On IBM i, end the standby by adding the `INSTANCE(*STANDBY)` option to the **ENDMQM** command.
- On Windows, UNIX, and Linux platforms, end the standby with the **endmqm -x QMgrName** command.

- b. If the queue manager is running as the active instance:


-  On IBM i, end the instance and transfer control to the standby instance by adding the `ALWSWITCH(*YES)` option to the **ENDMQM** command.  
If there is no standby instance running, the command fails, and you must start a standby instance on a different server.
- On Windows, UNIX, and Linux platforms, end the instance and transfer control to the standby instance with the **endmqm** command. For example, **endmqm -shutdown\_option -s QMgrName**, where *-shutdown\_option* is an optional parameter specifying the type of shutdown. For more information, see **endmqm**.

If there is no standby instance running, the command fails, and you must start a standby instance on a different server.

- c. If a queue manager is running as a single instance queue manager, you have no alternative but to stop the queue manager before applying the maintenance update.


When you complete this step, no queue manager instances are left running on the server you intend to update.

2. Continue with the maintenance update procedure, following the step to issue the **endmqm** command, or quiesce IBM MQ and apply maintenance to the IBM MQ server.
3. When you have completed the maintenance update, restart all the queue managers on the IBM MQ server, permitting standby instances:

-  On IBM i, add the `STANDBY(*YES)` option to the **STRMQM** command
- On Windows, UNIX, and Linux platforms, use the **strmqm -x QmgrName** command.

4. Repeat the procedure on the standby server, to update its maintenance level.

5. If necessary, switch the active instances back to the primary servers:

-  On IBM i, use the **ENDMQM** command with the `ALWSWITCH(*YES)` option, and the restart the instances using the **STRMQM** command with the `STANDBY(*YES)` option.
- On Windows, UNIX, and Linux platforms, use the **endmqm -shutdown\_option -s QMgrName** command, and the restart the instances using the **strmqm -x QmgrName** command.

### Related tasks:

“Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623  
When you apply and remove maintenance level updates to IBM MQ, no migration is required.  
Maintenance level updates are applied either as a fix pack, or by manually applying an interim fix.

## Migrating queue managers to new-function fix packs

This scenario illustrates running different levels of queue manager from a single installation using new-function fix packs. It contrasts migrating a queue manager to new command levels in new-function fix packs, to migrating a queue manager to a new command level in a new release. The scenario explains the relationship between new-function fix packs and maintenance fix packs.

**z/OS** New function fix-packs are not available on z/OS. On z/OS, the same result is achieved by using the **OPMODE** parameter to control queue manager migration; see “z/OS: OPMODE” on page 677.

### Before you begin

In this section, IBM WebSphere MQ Version 7.1 is used as the current release, and the release is denoted by *r* ; the subsequent release is denoted by *R*.

The scenario starts with a single installation of IBM WebSphere MQ Version 7.1, *Inst\_1*. *Inst\_1* is the primary installation; see Figure 72 on page 659. For illustration, there are two queue managers, *QM1* and *QM2*. *QM1* stays at the *7r0* command level, *QM2* moves to the highest command level available.

The use of version numbers and command levels is illustrative, and does not imply anything about future releases.

### About this task

Figure 71 on page 656 has time advancing down the Y-Axis, as new fix packs are released. On the X-Axis are different command levels. As a queue manager is migrated to a new command level, it shifts across the diagram. Each column represents the fix levels a queue manager at a particular command level can run at.

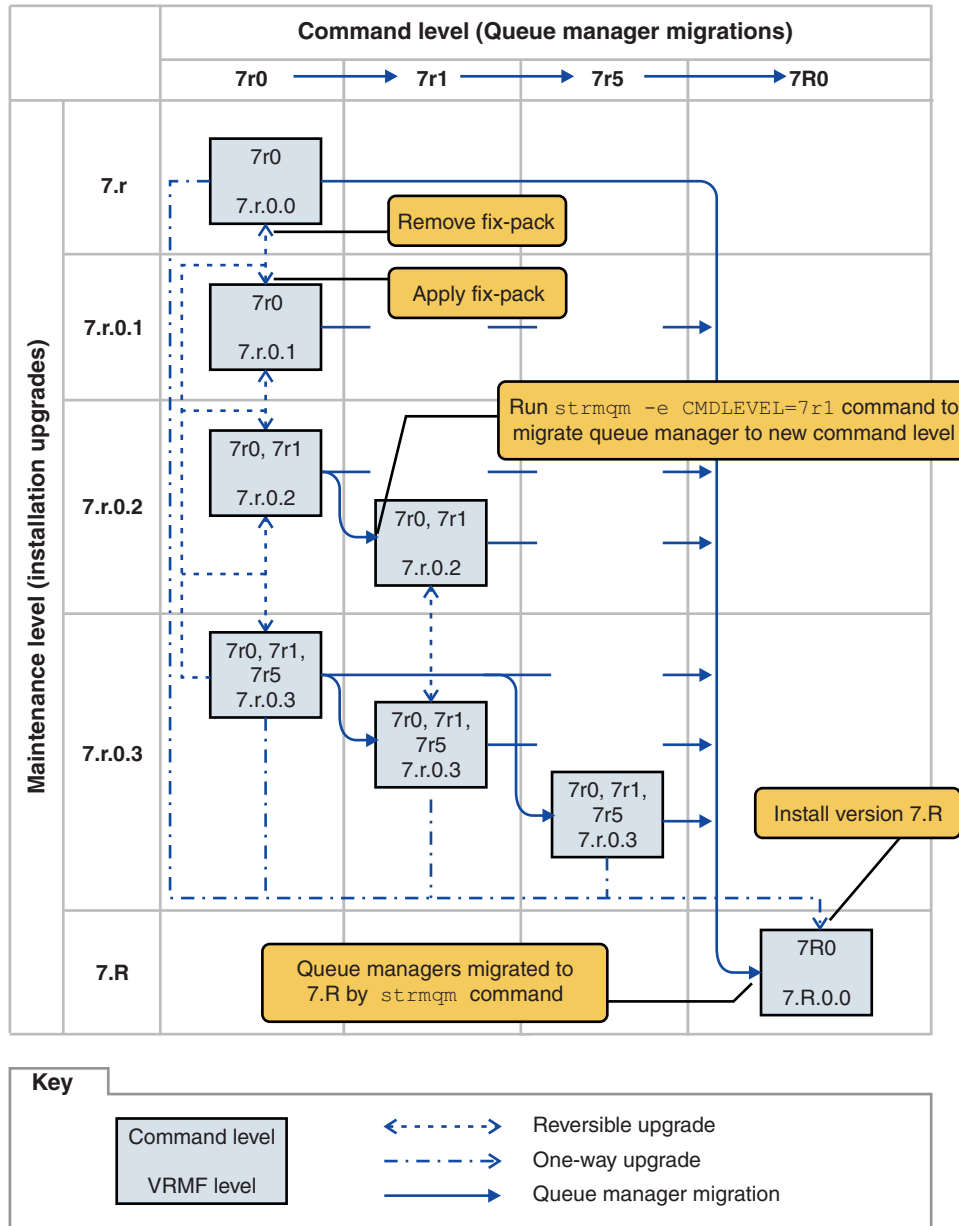


Figure 71. Migration of queue managers to new command levels using new-function fix packs

Figure 71 is a little complicated, but it captures lots of details about new-function fix packs to help you remember them. The steps in the task explain the details in the figure. Some of the principle features of Figure 71 are explained in the following list:

### Maintenance level and Command level

The maintenance level is a fix pack with a V.R.M.F. code; see “The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423. V.R.M.F codes are one to four digits, always punctuated by periods. Trailing zeros are sometimes omitted in descriptions, but never when a V.R.M.F code is used to label a fix pack. Version 8.0 is an example of using a V.R.M.F code to describe the version of IBM MQ.

The command level is the command level property of a queue manager; see `CommandLevel` (MQLONG). Command levels are three-digit codes.



Command levels and versions are related. Up to Version 7.1 the command level and the first three digits of the V.R.M.F. code always matched. From Version 7.1, with the introduction of new-function fix packs, the command level of a queue manager can be greater than the first three digits of an installation. The difference arises, if the queue manager has been associated with a new command level using the **strmqm** command.

From Version 7.1 the rule that links command levels and V.R.M.F levels has been changed. The rule is that when a new version of IBM MQ is released, it has a command level greater than released in a new-function fix pack in the previous release. Usually this means that a new release of IBM MQ changes the version or release level, rather than the maintenance level.

In Figure 71 on page 656, the maintenance level, on the Y-Axis is labeled with V.R.M.F codes, and the command level, on the X-Axis, with command levels. Note how the illustrative release of 7.R increases the released command level from 7r0 to 7R0, and 7R0 exceeds the highest command level shipped in a new-function fix pack, 7r5.

### Reversible and One-way upgrades

The mechanism to apply and remove fix packs varies by platform. You can apply any fix pack that changes only the maintenance or fix level of a release to an installation. Fix pack application is reversible. When you remove a fix pack, you restore the previous release level. So applying 7.r.0.3 to 7.r.0.1, and then removing it, restores the installation to 7.r.0.1.

Sometimes, you can change an installation to a particular V.R.M.F level by upgrading the installation with a “manufacturing refresh”. If you install a manufacturing refresh, you can only return to the earlier release level by uninstalling, and reinstalling; see “Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )” on page 453.

Applying a manufacturing refresh to modify the maintenance and fix level of a release is the same process as upgrading to a new version or release of IBM MQ. Neither can be reversed without uninstalling.

However there is a particular aspect of upgrading to a new version or release that is different from upgrading to a new maintenance or fix level. If you start a queue manager after a version or release upgrade, the command level of the queue manager is automatically increased. You can then no longer start the queue manager with the installation from the previous release.

On the diagram, an irreversible upgrade is shown by the “One-way” arrow between 7.r and 7.R. To prevent an accidental migration, you can rename the new installation. After renaming, rerun the **setmqm** command to associate a queue manager with the new release before running the **strmqm** command to migrate it.

If the upgrade applies only to the maintenance or fix level, you can restart the queue manager with the previous installation, if you reinstall it.

Manufacturing refresh maintenance releases are not distinguished from applying and removing fix packs on the diagram. Both are represented by reversible arrows in Figure 71 on page 656.

### Multiple installations

You might choose to have a different installation for each maximum command level supported by an installation. Each column on the diagram would represent a different installation.

You need only one installation at Version 7.1 to be able to select any command level released with Version 7.1 for a queue manager. Eventually, if you intend to run Version 7.1 and version 7.R in parallel, you must have two installations. The scenario that follows uses a single installation.

Another variation is to follow the “rolling fix pack” approach described in “UNIX, Linux, and Windows: Staging maintenance fixes” on page 669. You can maintain two installations at Version 7.1, one at the current fix level, and one either at a later or earlier fix level. You might then install version 7.R as a third installation, or replace the Version 7.1 installation at the older fix level.

### Migrating queue managers

The migration paths for queue managers are shown by solid arrows on the diagram. Some of the solid arrows are broken, to avoid cluttering the diagram with too many lines. If the migration to a higher command level jumps command levels, you do not have to migrate it through the intervening commands levels.

To migrate a queue manager to a higher command level in a new-function fix pack, you must start the queue manager with a special parameter:

```
▶▶ strmqm -e CMDLEVEL=Level QMGrName ▶▶
```

*Level* is the three-digit command level.

The queue manager stops immediately the migration process is complete. When you next start it, it runs at the new command level. The queue manager cannot be restarted at a lower command level. This rule means that you must associate the queue manager with an installation that includes a command level at least as great as the current command level of the queue manager.

### Restoring queue managers

To restore a queue manager to a lower command level, you must back up the queue manager before you migrate it to the higher command level.

### Procedure

This procedure keeps both QM1 and QM2 at the current maintenance level, QM1 at command level 7r0, and QM2 at the latest command level.

1. Download the next fix pack for the version of your product, for example, 7.1.0.2, when it is released. See Fix Central.

The initial system has two queue managers running 7.r.0.0 at command level 7r0 ; see Figure 72 on page 659.

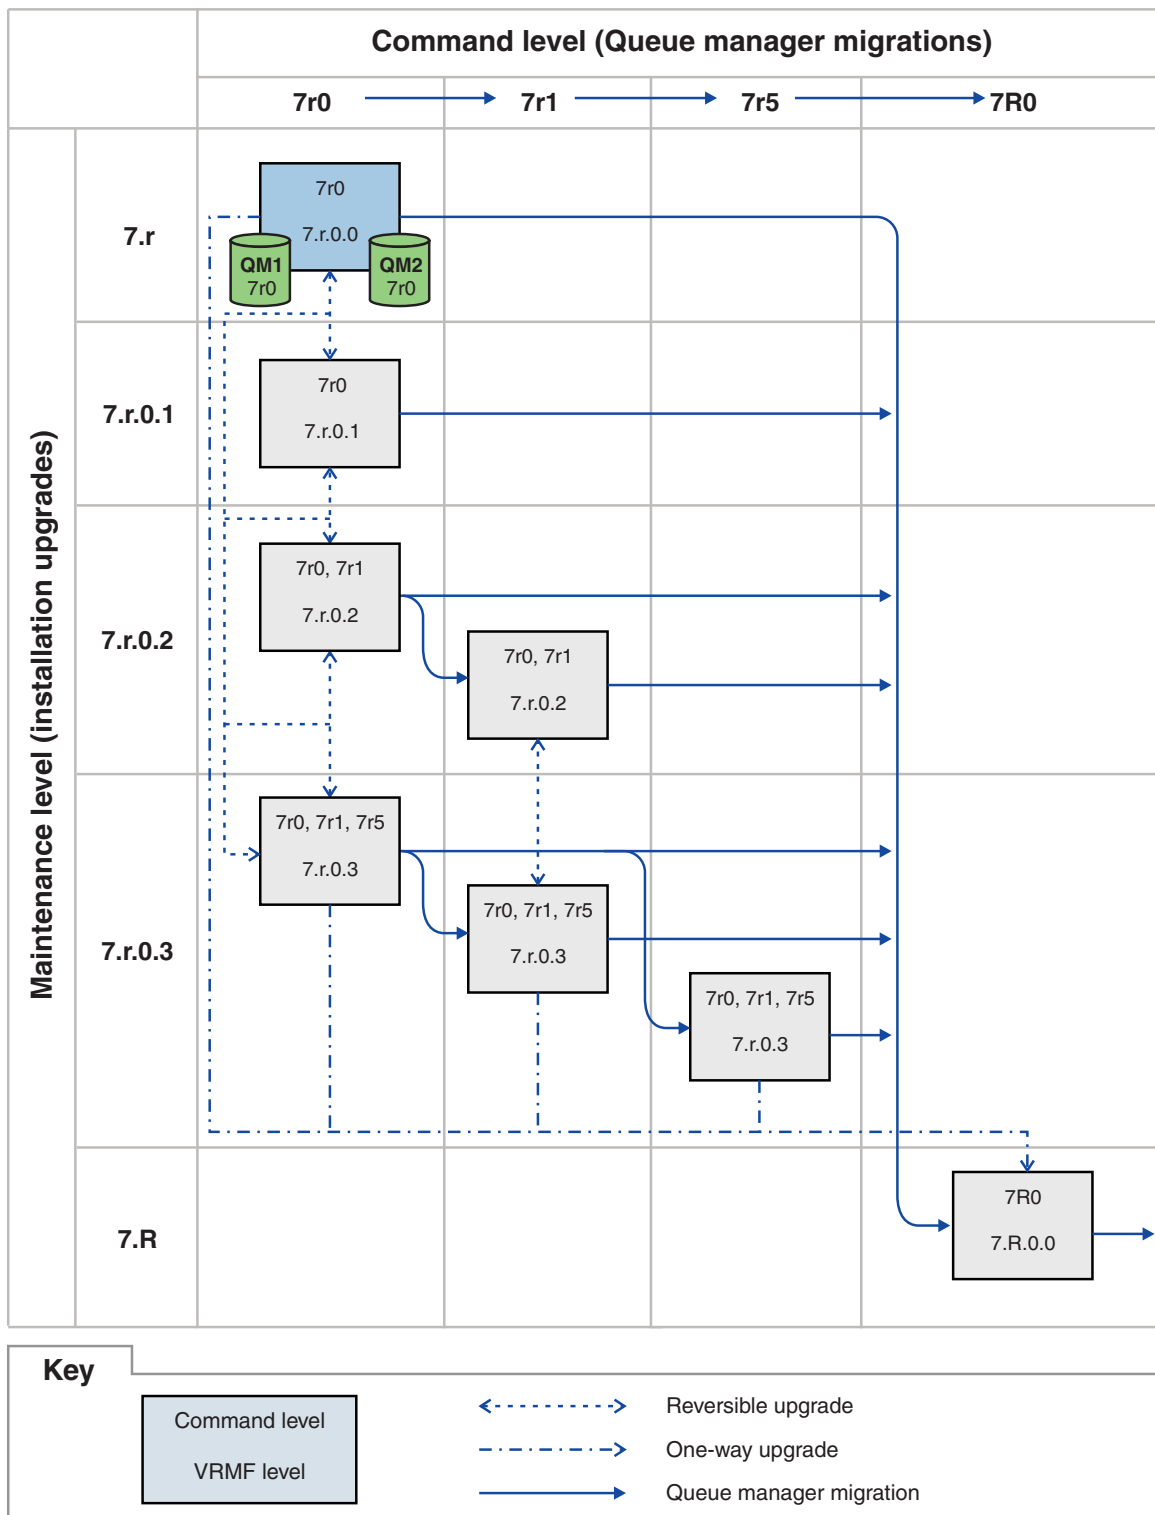


Figure 72. Initial state, QM1 and QM2 at command level 7r0, and fix level 7.r.0.0

2. Apply fix pack 7.r.0.1 to Inst\_1.

Follow the task for your platform in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.

3. Restart the queue managers.

Both queue managers are now running using Inst\_1 at the 7.r.0.1 maintenance level, and the 7r0 command level; see Figure 73 on page 661.

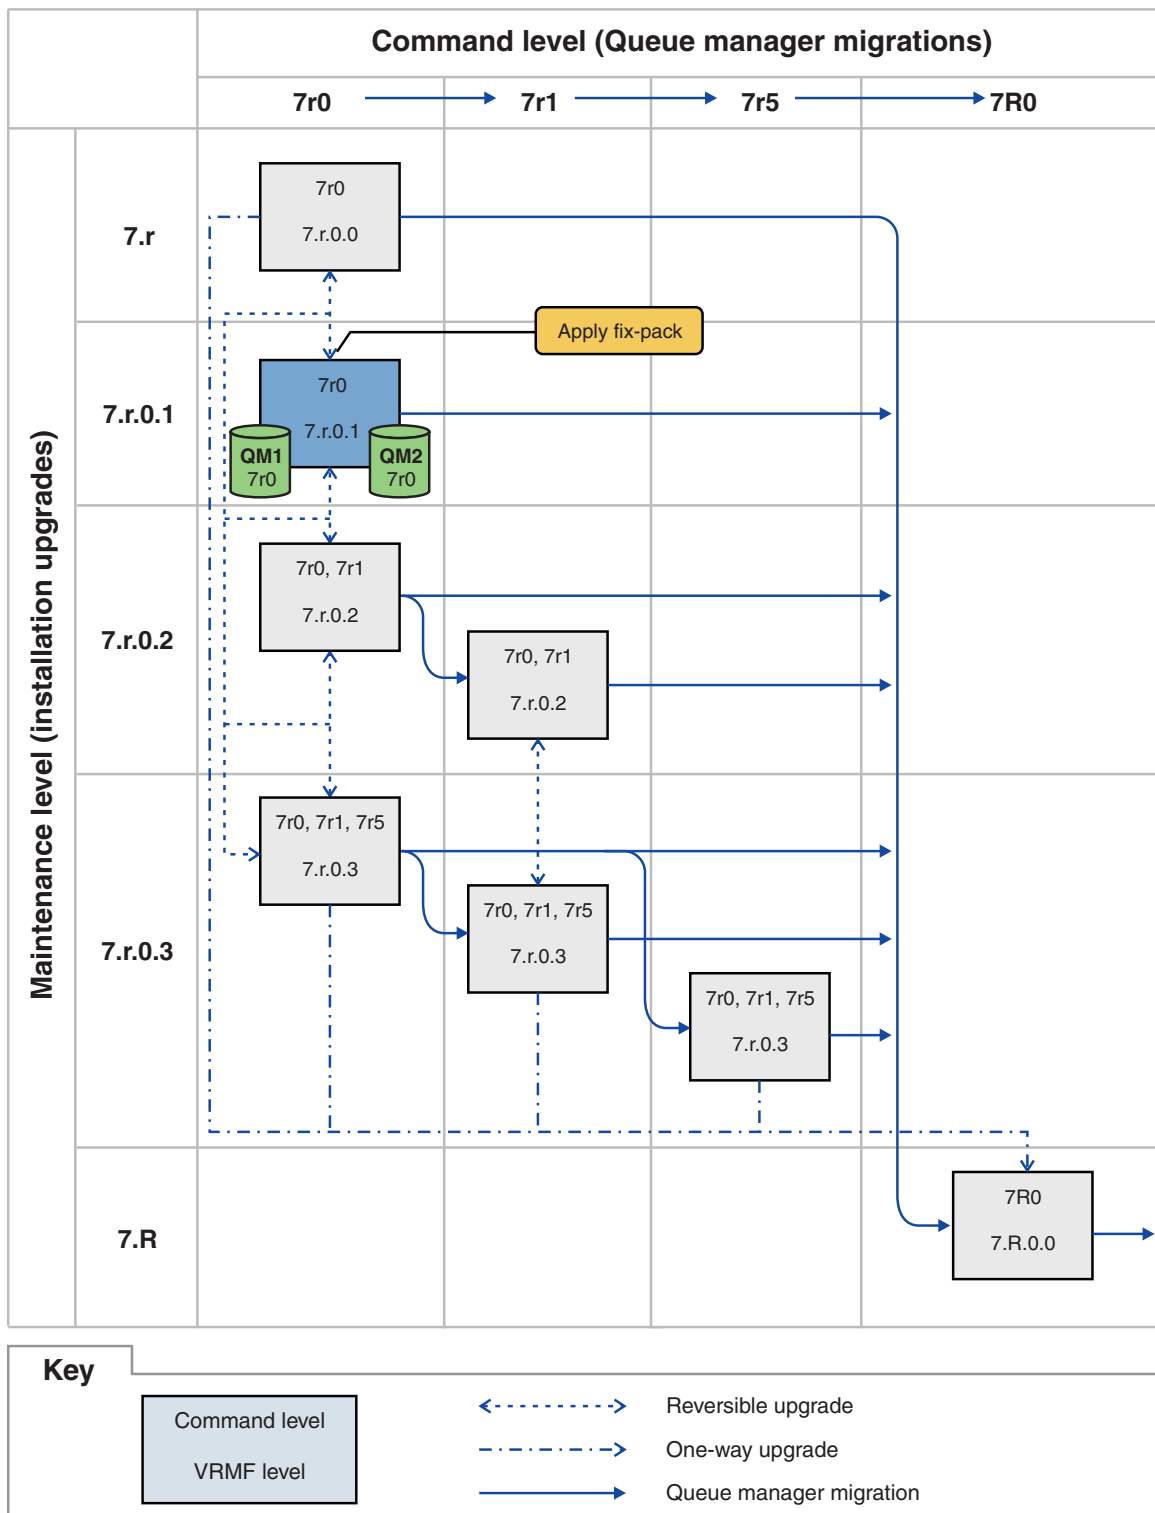


Figure 73. QM1 and QM2 at command level 7r0, and fix level 7.r.0.1

4. Apply fix pack 7.r.0.2.
  - a. Repeat steps 1 on page 658 and 2 on page 659 with fix pack 7.r.0.2.

5. Restart QM1.

QM1 is now running using Inst\_1 at the 7.r.0.2 maintenance level, and the 7r0 command level.  
The queue manager is not automatically migrated to the 7r1 command level.

6. Migrate QM2 to the 7r1 command level.

```
strmqm -e CMDLEVEL=711 QM2
```

QM2 is using Inst\_1 at the 7.r.0.2 maintenance level, and has been migrated to the 7r1 command level.

7. Restart QM2.

QM2 is now running using Inst\_1 at the 7.r.0.2 maintenance level, and the 7r1 command level;  
see Figure 74 on page 663.

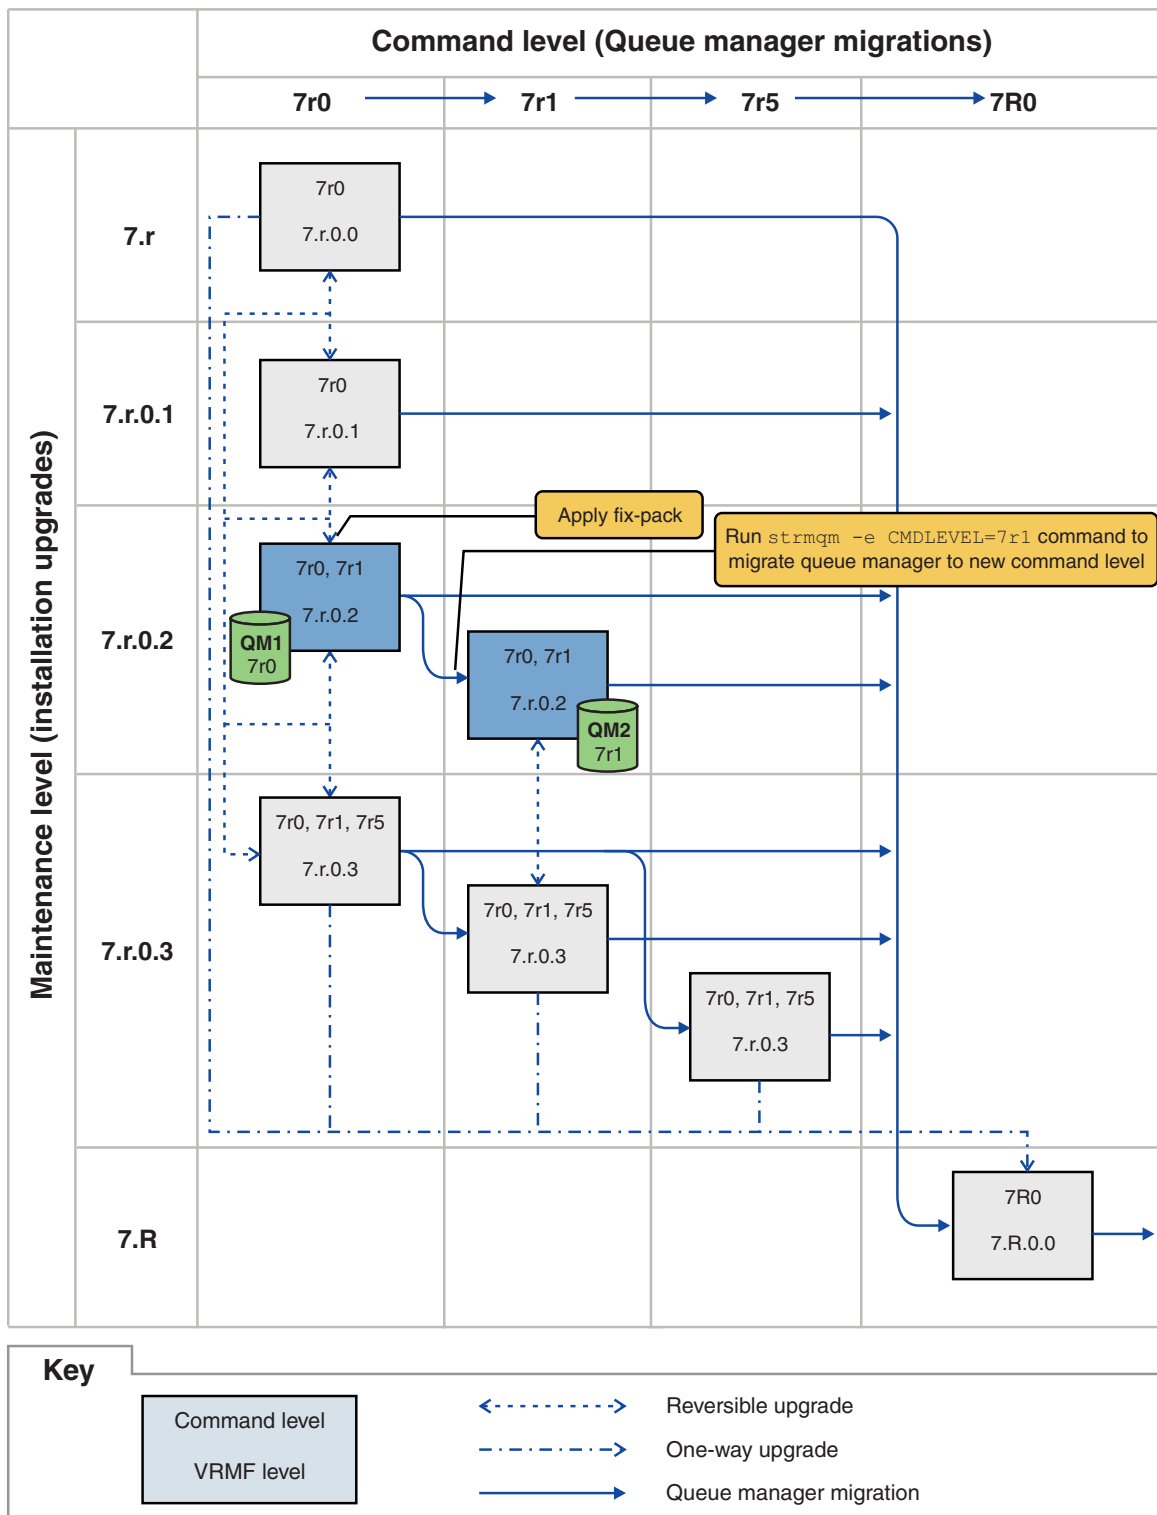


Figure 74. QM1 at command level 7r0 and fix level 7.r.0.2 ; QM2 at command level 7r1 and fix level 7.r.0.2

8. Apply fix pack 7.r.0.3 and migrate QM2 to the 7r5 command level.
  - a. Repeat steps 4 on page 661 to 5 on page 662 with fix pack 7.r.0.3.

- b. Repeat steps 6 on page 662 to 7 on page 662 with command level 7r5.
  - QM1 is using Inst\_1 at the 7.r.0.3 maintenance level, and is running at the 7r0 command level.
  - QM2 is using Inst\_1 at the 7.r.0.3 maintenance level, and has been migrated to the 7r5 command level; see Figure 75 on page 665.



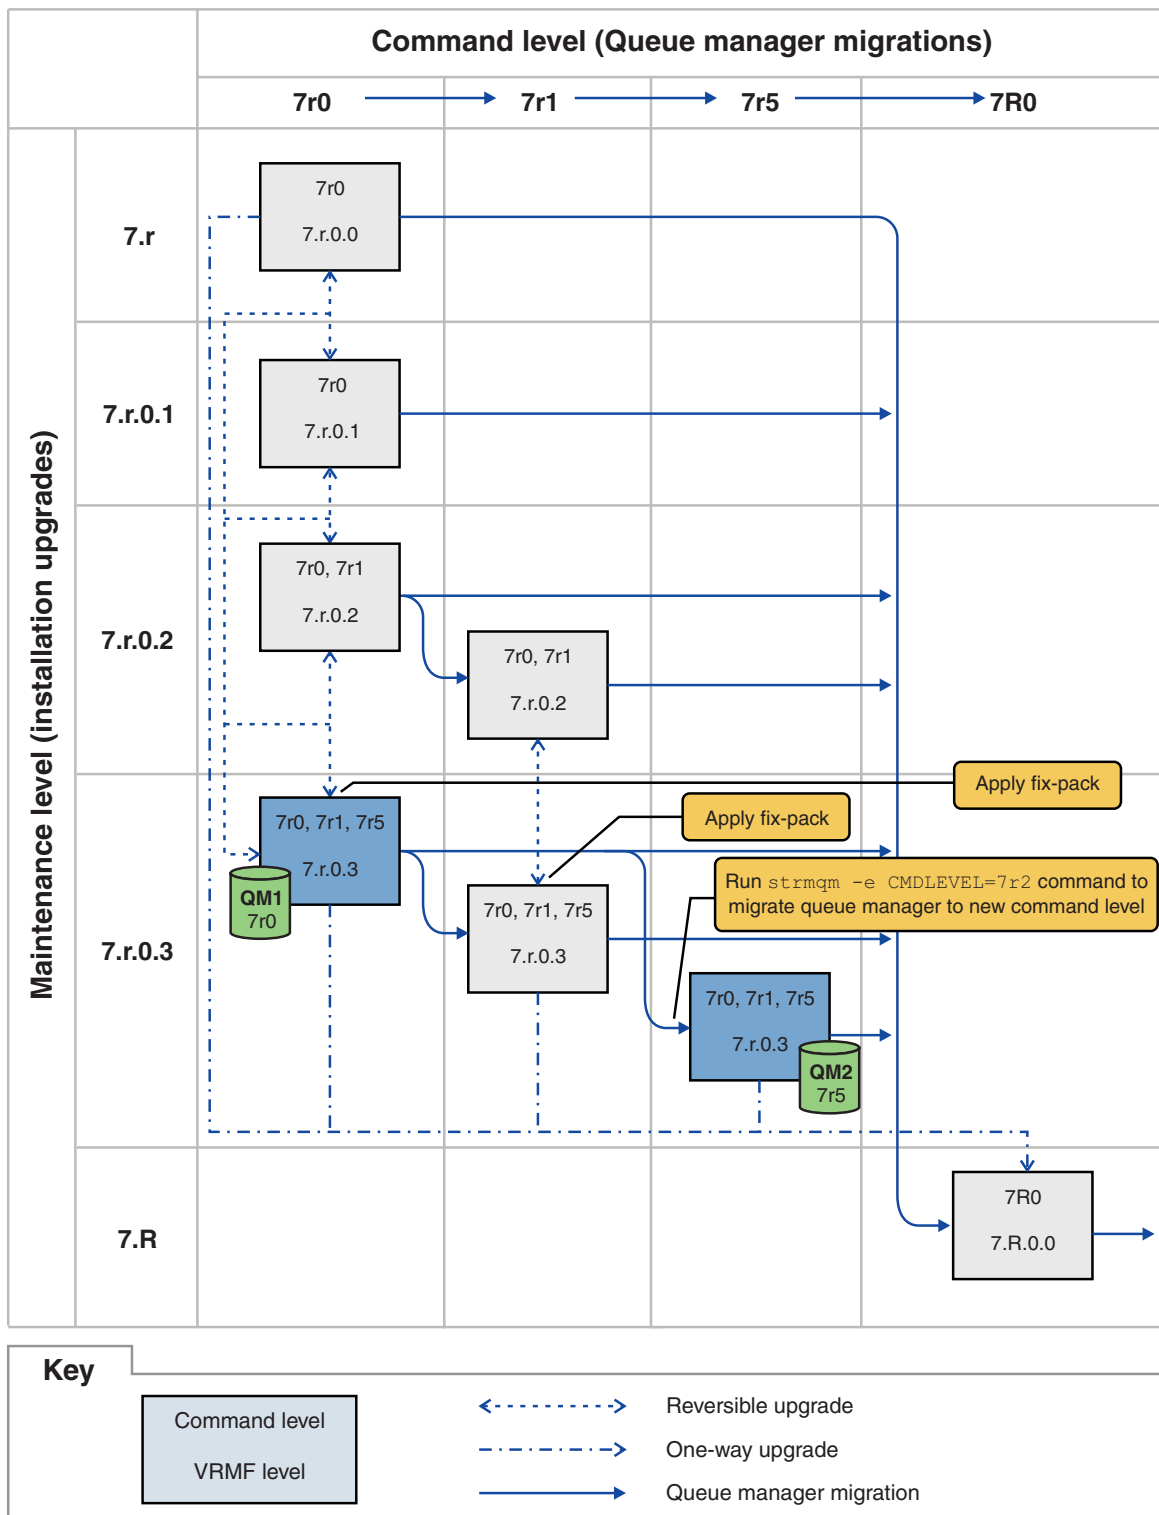


Figure 75. QM1 at command level 7r0 and fix level 7.r.0.3 ; QM2 at command level 7r5 and fix level 7.r.0.3

9. Migrate QM2 to version 7.R
  - On UNIX, Linux, and Windows:

- a. Install version 7.R, with the installation name `Inst_2`, alongside Version 7.1.
- b. Set up the local environment to the installation `Inst_2`.

Windows:

```
"Inst_2_INSTALLATION_PATH
\bin\setmqenv" -s
```

The `-s` option sets up the environment for the installation that runs the **setmqenv** command.

UNIX:

```
. Inst_2_INSTALLATION_PATH/bin/setmqenv -s
```

- c. Run the **setmqm** command to associate QM2 with `Inst_2`.

```
setmqm -m QM2 -n Inst_2
```

- d. Run the **strmqm** command to start QM2 and migrate it to version 7.R.

```
strmqm QM2
```

QM1 is using `Inst_1` at the 7.r.0.3 maintenance level, and is running at the 7r0 command level.

QM2 is using `Inst_2` at the 7.R.0.0 maintenance level, and has been migrated to the 7R0 command level; see Figure 75 on page 665.

`Inst_1` remains the primary installation.

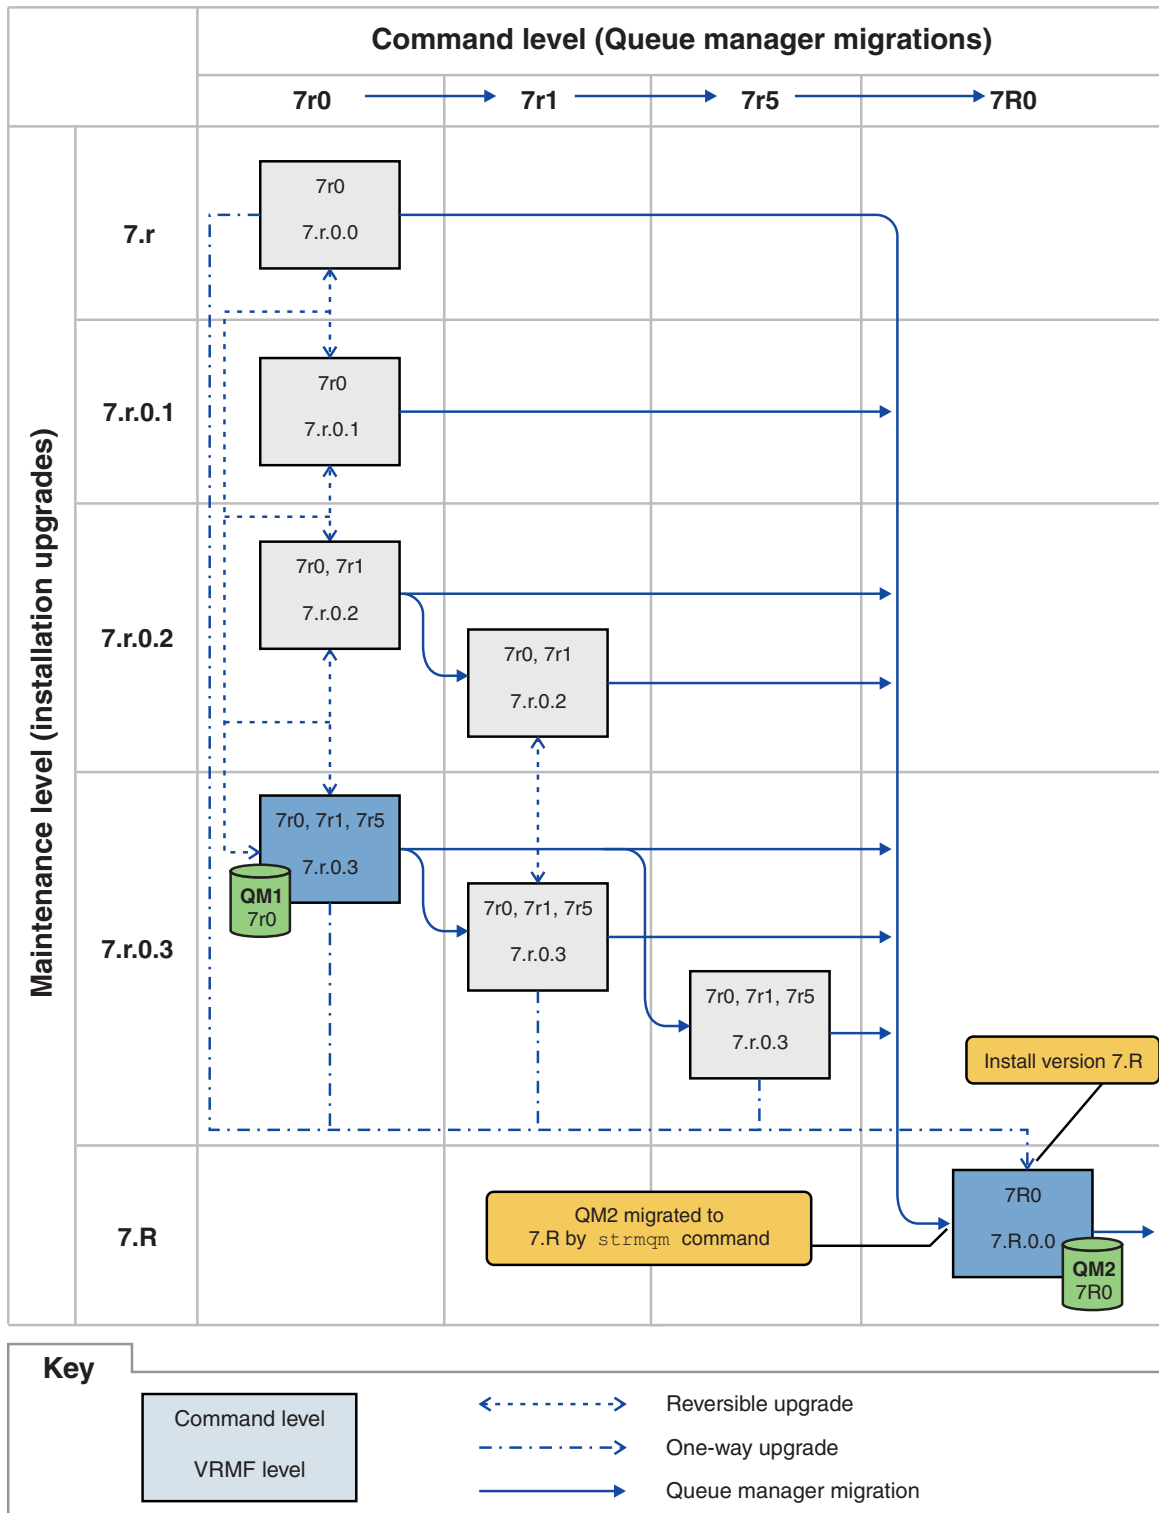



Figure 76. QM1 at command level 7r0 and fix level 7.r.0.3 ; QM2 at command level 7R0 and fix level 7.R.0.0


Related concepts:

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460  
You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.


“New function in maintenance level upgrades (On platforms other than z/OS )” on page 451  
IBM might introduce new functions between releases in maintenance level upgrades such as fix packs. A maintenance level upgrade including new function increases the maximum command level of an installation.

“Queue manager coexistence in Version 8.0” on page 457  
Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On  z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“The version naming scheme for IBM MQ (On platforms other than z/OS )” on page 423  
IBM MQ releases have a four-digit Version, Release, Maintenance, and Fix (VRMF) level code.

 “The version naming scheme for IBM MQ for z/OS” on page 422  
On IBM MQ for z/OS, releases have a three-digit Version, Release, and Maintenance (VRM) level code. The code is significant; it identifies the service life of a release. To run a queue manager at a different VRM level, you must migrate the queue manager, its applications, and the environment in which it runs. Depending on the migration path, the migration might require more or less effort.

“Upgrade, migration, and maintenance of IBM MQ (On platforms other than z/OS )” on page 453  
You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same UNIX, Linux, and Windows server. You can apply maintenance level upgrades to upgrade the maintenance or fix level. Applying maintenance level upgrades cannot change the version or release level of IBM MQ. Maintenance level upgrades can be reversed, installations cannot be reversed.

 “Upgrade, migration, and maintenance of IBM MQ on z/OS” on page 452  
You can install new releases of IBM MQ to upgrade IBM MQ to a new maintenance, release, or version level. Multiple installations at the same or different levels can coexist on the same z/OS instance. Running a queue manager at a higher level requires migration. Maintenance differs from upgrading. To maintain a level of IBM MQ, you apply Program Temporary Fixes (PTFs) to the installed code.

“IBM MQ maintenance” on page 441  
Maintenance is the application of a reversible fix. Any changes to queue manager data are compatible with the previous code level.

“IBM MQ migration” on page 448  
Migration is the conversion of programs and data to work with a new code level of IBM MQ. Some types of migration are required, and some are optional. Queue manager migration is never required after applying a maintenance level update, that does not change the command level. Some types of migration are automatic, and some are manual. Queue manager migration is typically automatic and required after releases and manual and optional after a maintenance level upgrade that introduces a new function. Application migration is typically manual and optional.

“IBM MQ upgrades and fixes” on page 442  
The term upgrade applies to changing the version V, release R, or modification M of a product. The term fix applies to a change in the F digit.

**Related tasks:**

“Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623  
When you apply and remove maintenance level updates to IBM MQ, no migration is required. Maintenance level updates are applied either as a fix pack, or by manually applying an interim fix.

“UNIX, Linux, and Windows: Staging maintenance fixes” on page 669  
Use multiple installations of IBM MQ on the same server to control the release of maintenance fixes.

**Related reference:**

z/OS “z/OS: OPMODE” on page 677

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. To access V8.0 capabilities, change the value of **OPMODE** to **OPMODE=(NEWFUNC,800)**. To restrict the use of new capabilities, and retain the ability to revert the queue manager to its earlier level, leave **OPMODE** at its default setting, **OPMODE=(COMPAT,800)**.

**Related information:**

CommandLevel (MQLONG)

## Querying the maintenance level

Query the IBM MQ maintenance level by running the **dspmqver** command

### About this task

After an update to the initial installation, the version indicates the maintenance level to which the product has been updated. For example, before applying any maintenance, the version is 7.0.1.0. As maintenance is applied the last digit is updated, for example to 7.0.1.3.

### Procedure

To view the version use the **dspmqver** command. At a command prompt, enter the following command: **dspmqver**. The resulting messages include the IBM MQ version number, which shows the maintenance level.

**Related information:**

**dspmqver**

## UNIX, Linux, and Windows: Staging maintenance fixes

Use multiple installations of IBM MQ on the same server to control the release of maintenance fixes.

### Before you begin

Set up your configuration modelled on the first row of Figure 77 on page 670. You can apply this scenario to any version of IBM MQ from IBM WebSphere MQ Version 7.1 onwards. In this scenario it is assumed you have a number of applications and two queue managers, QM1 and QM2, running on a server. IBM WebSphere MQ Version 7.0.1 is not installed on the server.

1. Install two copies of IBM MQ. In the example, they are named **Inst\_1** and **Inst\_2** and IBM WebSphere MQ Version 7.1 is being used.
2. Make **Inst\_1** primary by running **setmqinst**.
3. Associate all the queue managers on the server with **Inst\_1** by running **setmqm**.
4. Start all the queue managers on the server.
5. Show and connect all direct connections with the queue managers associated with **Inst\_1** in MQ Explorer.
6. Set up remote connections to all the queue managers in each instance of MQ Explorer.

### About this task

You can install multiple copies of IBM MQ on a server to stage the release of IBM MQ fixes. Figure 77 on page 670 illustrates a way of using two installations to roll out fixes. In this approach, you maintain two fix levels on a server, with the aim of getting all queue managers and applications to the production fix level before replacing the previous level on fix pack with the next level.

Which installation an application uses is driven by the queue manager an application connects to. The **setmqm** command associates a queue manager with an installation. You can associate a queue manager

with a different installation as long as the installation is at the same or higher command level. In this example, all the installations are at the same command level. You can associate or reassociate a queue manager with either of the installations running any of the fix packs.

In the example, an application links to the primary installation. When it connects to a queue manager, IBM MQ switches the linkage to the installation associated with the queue manager; see “Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460.

For applications built with the link options described in the product documentation, the simplest way to configure the link library search path for IBM MQ applications is to make an installation primary. Only if it is important to pick up a fix in the IBM MQ link library itself, must you review the search path. Either you must make the installation with the IBM MQ link library fix primary, or make a local adjustment for the application, perhaps by running the **setmqenv** command.

Running commands is a different matter. Commands are always run from the primary installation, or the installation you have selected by running the **setmqenv** command. If you run a command from the wrong installation, the command fails. For example, if QM1 is associated with Inst\_1, running the Windows command, Inst\_2\_Installation\_path/bin/strmqm QM1 fails.

If you are using MQ Explorer and you have two installations, you also have two MQ Explorer instances. One linked to one installation, and one to the other. Each MQ Explorer shows locally connected queue managers that are associated with the same installation as the instance of MQ Explorer. To monitor all the queue managers on a server, set up remote connections to the queue managers associated with the other installations.

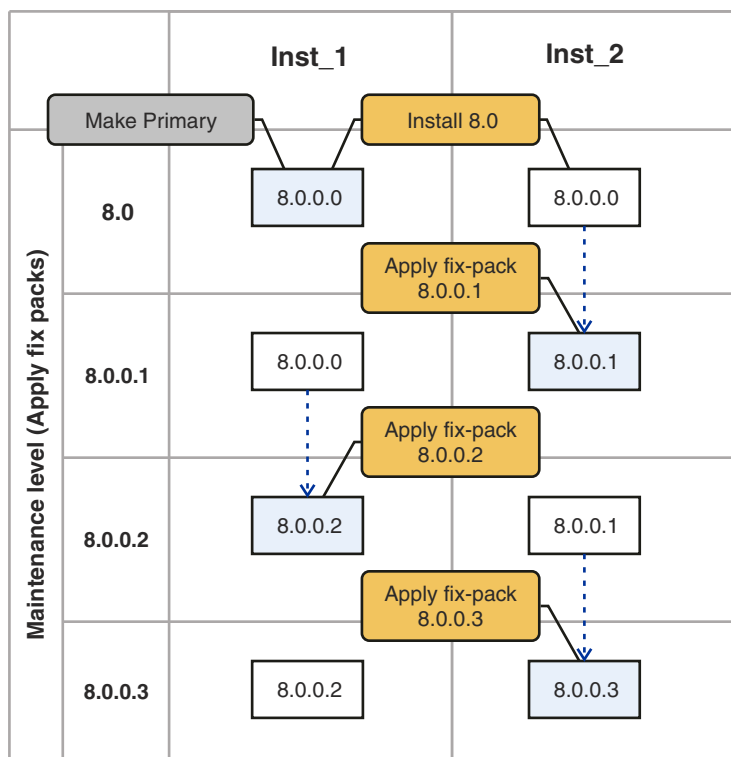


Figure 77. Rolling fix releases

## Procedure

1. Download the first fix pack, for example, 7.1.0.1, when it is released.  
See Fix Central.
2. Apply the fix pack you downloaded to Inst\_2.  
Follow the task for your platform in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.
3. Verify Inst\_2.
4. Transfer the queue managers to Inst\_2 one at a time.
  - a. Stop QM1 and the applications connected to it.  
The **endmqm** command informs an application that the queue manager it is connected to is stopping; see Stopping a queue manager.

**Note:** The topic, “Applying maintenance level upgrades to multi-instance queue managers” on page 653, describes how to apply maintenance to a multi-instance queue manager. A multi-instance queue manager can continue to run on one server, while maintenance is applied to another server.

- b. Set up the local environment to the installation Inst\_2.

**Windows** Windows:

```
"Inst_2_INSTALLATION_PATH\bin\setmqenv" -s
```

The -s option sets up the environment for the installation that runs the **setmqenv** command.

**UNIX** UNIX:

```
.Inst_2_INSTALLATION_PATH/bin/setmqenv -s
```

- c. Associate the queue manager with Inst\_2.

```
setmqm -m QM1 -n Inst_2
```

- d. Start QM1

```
strmqm QM1
```

- e. Repeat substeps c and d for QM2.
- f. Set up MQ Explorer for Inst\_2.
  - 1) Start the Inst\_2 instance of MQ Explorer

**Tip:** On Windows, hover over the IBM MQ icon in the system tray. The hover help shows the installation name associated with the MQ Explorer instance.

- 2) Click **IBM MQ > Queue Managers > Show/Hide Queue Managers... >**
- 3) Click each directly connected queue manager listed in the Hidden Queue Managers list > **Show**.
- 4) Click **Close**.

5. Set Inst\_2 primary.

**Windows** Windows:

```
"Inst_2_INSTALLATION_PATH\bin\setmqinst" -i -n Inst_2
```

**UNIX** UNIX:

```
Inst_2_INSTALLATION_PATH/bin/setmqinst -i -n Inst_2
```

6. Download the next fix pack for the version of your product, for example, 7.1.0.2, when it is released.  
See Fix Central.
7. Apply the fix pack that you have just downloaded to Inst\_1.  
Follow the task for your platform in “Applying and removing maintenance level updates (On platforms other than z/OS )” on page 623.
8. Verify Inst\_1.
9. Transfer queue managers to Inst\_1 one at a time.
  - a. Follow the procedure in step 4 on page 671  
Replacing Inst\_2 by Inst\_1 in the instructions.
10. Set Inst\_1 primary.

**Windows** Windows:

```
"Inst_1_INSTALLATION_PATH\bin\setmqinst" -i -n Inst_1
```

**UNIX** UNIX:

```
Inst_1_INSTALLATION_PATH/bin/setmqinst -i -n Inst_1
```

11. Repeat steps 1 on page 671 to 5 on page 671 for the odd-numbered fix packs of your product.
12. Repeat steps 6 to 10 for the even-numbered fix packs of your product.

#### **Related concepts:**

“Queue manager coexistence in Version 8.0” on page 457

Queue managers, with different names, can coexist on any server as long as they use the same IBM MQ installation. On **z/OS** z/OS, UNIX, Linux, and Windows, different queue managers can coexist on the same server and be associated with different installations.

“Multi-installation queue manager coexistence on UNIX, Linux, and Windows” on page 460

You can install multiple copies of IBM MQ for UNIX, Linux, and Windows on the same server. The installations must be at Version 7.1 or later, with one exception. One Version 7.0.1 installation, at fix pack level 6, or later, can coexist with multiple Version 7.1, or later installations.

#### **Related tasks:**

“Migrating IBM MQ library loading from an earlier version of the product to the latest version” on page 533

No change in the way IBM MQ libraries are loaded is normally required if you upgrade from an earlier version of the product to the latest version. You must have followed the instructions on building IBM MQ applications in Version 7.0.1 and you must replace IBM WebSphere MQ Version 7.0.1 with the latest version of the product. If you choose to take advantage of multi-installation in the latest version of the product, based on the side-by-side or multi-stage migration scenarios, you must modify the environment for the operating system to resolve IBM MQ dependencies for an application. Typically, you can modify the runtime environment, rather than relink the application.

#### **Related information:**

Installing IBM MQ server on Windows

Associating a queue manager with an installation

Changing the primary installation



setmqenv  
setmqinst  
setmqm

---

## Migration commands, utilities, and reference information

A selection of commands, utilities, and application reference information related to migration are collected together in the following subtopics.

### JMS PROVIDERVERSION property

The JMS **PROVIDERVERSION** property selects whether a Java application publishes and subscribes using the queued command message interface, or the integrated call interface.

#### Related information:

Configuring the JMS **PROVIDERVERSION** property

### Rules for selecting the IBM MQ messaging provider mode

The IBM MQ messaging provider has three modes of operation: normal mode, normal mode with restrictions, and migration mode. You can select which of these modes of operation a JMS application uses to publish and subscribe by setting the **PROVIDERVERSION** property for the connection factory to the appropriate value. In some cases, the **PROVIDERVERSION** property is set as unspecified, in which case an algorithm is used to determine which mode of operation to use.

If you cannot change the connection factory that you are using, you can use the `com.ibm.msg.client.wmq.overrideProviderVersion` property to override any setting on the connection factory. This override applies to all connection factories in the JVM but the actual connection factory objects are not modified.

You can set the **PROVIDERVERSION** property to any of the values 8 (normal mode), 7 (normal mode with restrictions), 6 (migration mode), or unspecified (the default value). The value that you specify for the **PROVIDERVERSION** property must be a string. If you are specifying an option of 8, 7 or 6, you can do this in any of the following formats:

- V.R.M.F
- V.R.M
- V.R
- V

where V, R, M and F are integer values greater than or equal to zero. The extra R, M and F values are optional and are available for you to use in case fine grained control is needed. For example, if you wanted to use a **PROVIDERVERSION** level of 7, you could set **PROVIDERVERSION=7, 7.0, 7.0.0 or 7.0.0.0**.

#### 8 - Normal mode

The JMS application uses the IBM MQ messaging provider normal mode. Normal mode uses all the features of a IBM MQ queue manager to implement JMS. This mode is optimized to use the JMS 2.0 API and functionality.

If you are connecting to a queue manager with a command level of 800, then all of the JMS 2.0 API and features, such as asynchronous send, delayed delivery, or shared subscription, can be used.

If the queue manager specified in the connection factory settings is not a Version 8.0.0 queue manager, the `createConnection` method fails with an exception `JMSFMQ0003`.

The IBM MQ messaging provider normal mode uses the sharing conversations feature and the number of conversations that can be shared is controlled by the **SHARECNV()** property on the server connection channel. If this property is set to 0, you cannot use IBM MQ messaging provider normal mode and the `createConnection` method fails with an exception `JMSCC5007`.

## 7 - Normal mode with restrictions

The JMS application uses the IBM MQ messaging provider normal mode with restrictions. This mode uses the JMS 2.0 API, but not the new features such as shared subscriptions, delayed delivery, or asynchronous send.

If you set **PROVIDERVERSION** to 7 only the IBM MQ messaging provider normal with restrictions mode of operation is available. If the queue manager specified in the connection factory settings is not a Version 7.0.1, or later, queue manager, the `createConnection` method fails with exception `JMSFCC5008`.

If you are connecting using normal mode with restrictions, to a queue manager with a command level between 700 and 800 then you can use the JMS 2.0 API, but not the asynchronous send, delayed delivery, or shared subscription features.

The IBM MQ messaging provider normal mode with restrictions uses the sharing conversations feature and the number of conversations that can be shared is controlled by the **SHARECNV()** property on the server connection channel. If this property is set to 0, you cannot use IBM MQ messaging provider normal mode with restrictions and the `createConnection` method fails with an exception `JMSCC5007`.

## 6 - Migration mode

The JMS application uses the IBM MQ messaging provider migration mode.

The IBM MQ classes for JMS use the features and algorithms supplied with IBM WebSphere MQ Version 6.0. If you want to connect to WebSphere Message Broker Version 6.0 or 6.1 using IBM WebSphere MQ Enterprise Transport Version 6.0, you must use this mode. You can connect to a IBM MQ Version 8.0 queue manager using this mode, but none of the new features of a IBM MQ classes for JMS queue manager are used, for example, read ahead or streaming.

If you have a IBM MQ Version 8.0 client connecting to a IBM MQ Version 8.0 queue manager, then the message selection is done by the queue manager rather than on the client system.

If IBM MQ messaging provider migration mode is specified and you attempt to use any of the JMS 2.0 API, the API method call fails with the exception `JMSCC5007`.

### unspecified (default)

The **PROVIDERVERSION** property is set to *unspecified* by default.

A connection factory that was created with a previous version of IBM MQ classes for JMS in JNDI takes this value when the connection factory is used with the new version of IBM MQ classes for JMS. The following algorithm is used to determine which mode of operation is used. This algorithm is used when the `createConnection` method is called and uses other aspects of the connection factory to determine if IBM MQ messaging provider normal mode, normal mode with restrictions, or IBM MQ messaging provider migration mode is required.

1. First, an attempt to use IBM MQ messaging provider normal mode is made.
2. If the queue manager connected is not IBM MQ Version 8.0, or later, an attempt to use IBM MQ messaging provider normal mode with restrictions is made.
3. If the queue manager connected is not IBM WebSphere MQ Version 7.0.1, or later, the connection is closed and IBM MQ messaging provider migration mode is used instead.
4. If the **SHARECNV** property on the server connection channel is set to 0, the connection is closed and IBM MQ messaging provider migration mode is used instead.
5. If **BROKERVER** is set to V1 or the default *unspecified* value, IBM MQ messaging provider normal mode continues to be used, and therefore any publish/subscribe operations use the new IBM WebSphere MQ Version 7.0.1, or later, features.

See `ALTER QMGR` for information about the `PSMODE` parameter of the `ALTER QMGR` command for further information on compatibility.

6. If **BROKERVER** is set to V2 the action taken depends on the value of **BROKERQMGR** :
  - If the **BROKERQMGR** is blank:

If the queue specified by the **BROKERCONQ** property can be opened for output (that is, MQOPEN for output succeeds) and **PSMODE** on the queue manager is set to COMPAT or DISABLED, then IBM MQ messaging provider migration mode is used.

- If the queue specified by the **BROKERCONQ** property cannot be opened for output, or the **PSMODE** attribute is set to ENABLED:  
IBM MQ messaging provider normal mode is used.
- If **BROKERQMGR** is non-blank :  
IBM MQ messaging provider migration mode is used.

**Related concepts:**

“When to use **PROVIDERVERSION**”

In two cases you must override the default selection of **PROVIDERVERSION** for the IBM MQ classes for JMS to work correctly.

**Related information:**

BROKERQMGR

BROKERCONQ

PSMODE

JMS provider version troubleshooting

**When to use PROVIDERVERSION**

In two cases you must override the default selection of **PROVIDERVERSION** for the IBM MQ classes for JMS to work correctly.

**Note:**

The migration mode that is described in this topic is migration from IBM WebSphere MQ Version 6.0 to IBM WebSphere MQ Version 7.0.

IBM WebSphere MQ Version 6.0, WebSphere Application Server Version 6.0.x, and WebSphere Message Broker Version 6 are out of support, and therefore this topic is included only for reference purposes.

There are two scenarios where you cannot use the algorithm described in “Rules for selecting the IBM MQ messaging provider mode” on page 673. Consider using **PROVIDERVERSION** in these scenarios.

1. If IBM Integration Bus is in compatibility mode, you must specify **PROVIDERVERSION** for it to work correctly.
2. If you are using WebSphere Application Server Version 6.0.1, WebSphere Application Server Version 6.0.2, or WebSphere Application Server Version 6.1, connection factories are defined using the WebSphere Application Server administrative console.

In WebSphere Application Server the default value of the **BROKERVER** property on a connection factory is V2. The default **BROKERVER** property for connection factories created by using **JMSAdmin** or IBM MQ Explorer is V1. This property is now “unspecified” in IBM MQ.

If **BROKERVER** is set to V2 (either because it was created by WebSphere Application Server or the connection factory has been used for publish/subscribe before) and has an existing queue manager that has a **BROKERCONQ** defined (because it has been used for publish/subscribe messaging before), the IBM MQ messaging provider migration mode is used.

However, if you want the application to use peer-to-peer communication and the application is using an existing queue manager that has ever done publish/subscribe, and has a connection factory with **BROKERVER** set to 2 (if the connection factory was created in WebSphere Application Server this is the default), the IBM MQ messaging provider migration mode is used. Using IBM MQ messaging provider migration mode in this case is unnecessary; use IBM MQ messaging provider normal mode instead. You can use one of the following methods to work around this:

- Set **BROKERVER** to 1 or unspecified. This is dependent on your application.
  - Set **PROVIDERVERSION** to 8, or 7 which are custom properties in WebSphere Application Server Version 6.1.
- Alternatively, use the client configuration property, or modify the queue manager connected so it does not have the **BROKERCONQ**, or make the queue unusable.

## strmqbrk: Migrate the IBM WebSphere MQ Version 6.0 publish/subscribe broker to a later version

Migrate the persistent state of an IBM MQ publish/subscribe broker to a later version queue manager.

### Purpose

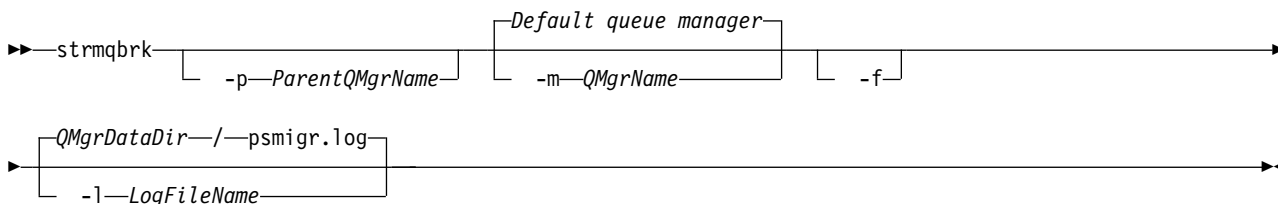
Use the **strmqbrk** command to migrate the state of a IBM WebSphere MQ Version 6.0 publish/subscribe broker to a later version queue manager. If the queue manager has already been migrated, no action is taken.

In IBM WebSphere MQ Version 6.0, **strmqbrk** started a broker. IBM MQ Version 8.0 publish/subscribe cannot be started in this manner. To enable publish/subscribe for a queue manager, use the **ALTER QMGR** command.

You can also use the **runmqbrk** command. This has the same parameters as **strmqbrk** and exactly the same effect.

### Syntax

AIX, HP-UX, Linux, Solaris, and Windows



### Optional parameters

AIX, HP-UX, Linux, Solaris, and Windows

**-p** *ParentQMGrName*

**Note:** This option is deprecated. **strmqbrk** migrates the parent connection automatically. If you specify the current parent queue manager, a warning message is issued and migration continues. If you specify a different queue manager, a error is issued and migration is not performed.

**-m** *QMGrName*

The name of the queue manager to be migrated. If you do not specify this parameter, the command is routed to the default queue manager.

**-f** Force migration. This option specifies that objects created during the migration replace existing objects with the same name. If this option is not specified, if migration would create a duplicate object, a warning is issued, the object is not created, and migration continues.

**-l** *LogFileName*

Log migration activity to the file specified in *LogFileName*.

## Syntax

IBM i

```
▶▶ STRMQMBRK ───▶▶
 ┌──┐ ┌──────────────────────────────────┐
 └─PARENTMQM──(ParentQMgrName)──┘ └─MQMNAME──QMgrName──┘
```

## Optional parameters

IBM i

**-PARENTMQM** *(ParentQMgrName)*

**Note:** This option is deprecated.

If you specify the current parent queue manager, a warning message is issued and migration continues. If you specify a different queue manager, a warning is issued and migration is not performed.

**-MQMNAME** *QMgrName*

The name of the queue manager to be migrated. If you do not specify this parameter, the command is routed to the default queue manager.

### Related information:

ALTER QMGR

## z/OS: OPMODE

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. To access V8.0 capabilities, change the value of **OPMODE** to **OPMODE=(NEWFUNC,800)**. To restrict the use of new capabilities, and retain the ability to revert the queue manager to its earlier level, leave **OPMODE** at its default setting, **OPMODE=(COMPAT,800)**.

The default setting of **OPMODE** at V8.0 is **OPMODE=(COMPAT,800)**. The default settings for a queue manager, when initially run at V8.0, restrict it to a limited set of the new functions. The restriction makes it possible to revert a queue manager to its earlier release level, if you must do so.

If you query the value of **OPMODE** with the command **DIS SYSTEM**, the result is **(NEWFUNC, VRM)**, where **VRM** is a release level. **VRM** is the same as the queue manager command level. The value of **VRM** is the release level you can revert the queue manager back to. If the queue manager is newly created at V8.0, then **VRM=800**. If the queue manager was previously run on V7.1.0, then **VRM=710**.

Make all the new functions in the release available to a queue manager by setting **OPMODE** to **OPMODE=(NEWFUNC,800)**. If you change the setting of **OPMODE** for a queue manager to **OPMODE=(NEWFUNC,800)**, you can no longer revert the queue manager to run on an earlier release level.

Each queue manager in a queue sharing group (QSG) must have **OPMODE** set to **OPMODE=(NEWFUNC,800)** and be restarted in order for any queue manager in the QSG to utilize Version 8.0 function.

This means that there are effectively two phases to enabling Version 8.0 new function in a QSG:

1. The first restart with **OPMODE** set to **OPMODE=(NEWFUNC,800)** prevents the queue manager from being reverted to run on an earlier release level.
2. The second restart, when all other QSG members have already been restarted with **OPMODE** set to **OPMODE=(NEWFUNC,800)**, and therefore cannot revert to run on an earlier release level, allows Version 8.0 new function to be used.

For example, in a QSG containing three queue managers, there will be a total of five queue manager restarts required to enable Version 8.0 new function on all QSG members.

You can reset **OPMODE** to `OPMODE=(COMPAT,800)`, after setting `OPMODE=(NEWFUNC,800)`, to prevent new functions being used. If you do so, **DIS SYSTEM** shows `OPMODE=(COMPAT,800)` rather than `OPMODE=(COMPAT,710)`, indicating that you cannot revert the queue manager to V7.1.0.

The syntax of **OPMODE** is as follows:

**OPMODE**=( *Mode*,*VerificationLevel* )

**OPMODE** specifies the operation mode of the queue manager.

The default setting of **OPMODE** is `OPMODE=(COMPAT,800)` .

#### *Mode*

Specifies the requested operation mode. The values are as follows:

#### **COMPAT**

The queue manager runs in compatibility mode. Certain new functions are not available. The queue manager can be migrated back to an earlier release.

#### **NEWFUNC**

All new functions provided in this level of code are available. The queue manager cannot be migrated back to an earlier release.

#### *VerificationLevel*

*VerificationLevel* is a Version.Release.Modification (VRM) code, without punctuation; 800, for example.

The value of *VerificationLevel* ensures that the **CSQ6SYSP** parameters are coded for use with the level of **CSQ6SYSP** macro being compiled. If *VerificationLevel* does not match the VRM level of SCSQMACS used for **CSQ6SYSP**, then a compile-time error is reported. The *VerificationLevel* is compiled into the parameter module.

At queue manager startup, if the *VerificationLevel* does not match the release level of the queue manager, then COMPAT mode is forced.

The intent of the *VerificationLevel* parameter is to avoid inadvertent and irreversible setting of **OPMODE** to NEWFUNC. The mistake might occur when migrating to a newer version of IBM MQ using **CSQ6SYSP** statements prepared for an older version of the queue manager. It might also occur using a **CSQ6SYSP** parameter module built with an older version of the SCSQMACS macros.

If you need assistance to revert to an earlier version of IBM MQ, contact your IBM support center.

#### **Related information:**

z/OS: Switching from `OPMODE=(NEWFUNC,800)` to `OPMODE=(COMPAT,800)`

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. You should be aware of the implications of switching from `OPMODE=(NEWFUNC,800)` to `OPMODE=(COMPAT,800)`

Using CSQ6SYSP

DISPLAY SYSTEM

#### **z/OS: OPMODE restrictions by version**

The availability of some new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. The **OPMODE** parameter determines whether you can use selected new functions, before you commit to staying at a given release. The functions and capabilities that are restricted in the different versions of the product are listed here.

#### **IBM WebSphere MQ Version 7.0.1**

Table 104. Functions and capabilities in Version 7.0.1

| Function                                               | Reference                              |
|--------------------------------------------------------|----------------------------------------|
| Log compression                                        | Log compression                        |
| GROUP units of recovery                                | Unit of recovery disposition           |
| Publish/subscribe migration using the CSQUMGMB utility | CSQUMGMB removed in IBM MQ Version 8.0 |

## IBM WebSphere MQ Version 7.1

Table 105. Functions and capabilities in Version 7.1

| Function                                                                                        | Reference                                                           |
|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| CFLEVEL(5) which is a prerequisite for Shared Message Data Sets (SMDS) and custom offload rules | DEFINE CFSTRUCT Planning your coupling facility and offload storage |
| Support for the X'3C' OTMA protocol messages, which report IMS health status                    | The IMS bridge                                                      |
| Resilience to coupling facility connectivity failures                                           | Shared queue recovery                                               |

## IBM MQ Version 8.0

Table 106. Functions and capabilities in Version 8.0

| Function                                                                                               | Reference                                                                                                                                                                                        |
|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Buffer pools can be located above the bar                                                              | DEFINE BUFFPOOL                                                                                                                                                                                  |
| Multiple cluster transmission queues                                                                   | Clustering: Planning how to configure cluster transmission queues Working with cluster transmission queues and cluster-sender channels Planning how you use multiple cluster transmission queues |
| Increased maximum addressable log range                                                                | Larger log Relative Byte Address                                                                                                                                                                 |
| Configurable certificates for individual channels using CERTLABL                                       | Certificate label (CERTLABL)                                                                                                                                                                     |
| Using host names, restricting the certificate issuer, and checking client credentials in CHLAUTH rules | SET CHLAUTH                                                                                                                                                                                      |

### Related information:

Using CSQ6SYSP

## z/OS: Switching from OPMODE=(NEWFUNC,800) to OPMODE=(COMPAT,800)

The availability of new functions and backward migration for IBM MQ for z/OS is controlled by the **OPMODE** parameter in the **CSQ6SYSP** macro. You should be aware of the implications of switching from **OPMODE=(NEWFUNC,800)** to **OPMODE=(COMPAT,800)**

### Switching from OPMODE=(NEWFUNC,800) to OPMODE=(COMPAT,800)

When you switch from **OPMODE=(NEWFUNC,800)** to **OPMODE=(COMPAT,800)** the following conditions occur:

- If the BSDS has been converted to version 2, the queue manager will not be able to access the BSDS when it is started at **OPMODE=(COMPAT,800)**. This means that the queue manager fails to start and terminates with reason code 00D10120.
  - You can find the BSDS version by running the print log map utility (CSQJU004).

- Any buffer pools with an ID greater than 15 are marked as suspended. This means that these buffer pools cannot be used, deleted, or altered until OPMODE=(NEWFUNC,800) is specified again. Information about the buffer pools is kept in checkpoint log records until OPMODE=(NEWFUNC,800) is specified again.
  - Any page set that uses a suspended buffer pool is also suspended. Information about the suspended page set is also kept in checkpoint records.
  - While a page set is suspended, any messages in the page set are unavailable. An attempt to use a queue or topic which uses the suspended page set results in an MQRC\_PAGESET\_ERROR message.
  - While it is suspended, a page set can be associated with a different buffer pool by using the FORMAT function of the utility program CSQUTIL, specifying TYPE(REPLACE). You can then issue a **DEFINE PSID** command to bring the page set back into use with a different buffer pool.

**Note:** All units of recovery that involved the suspended page set, except units that are indoubt, will have been backed out by the queue manager when the page set was last used. Indoubt units of recovery can be resolved when the page set is again in use by the queue manager.

- Any buffer pools with an ID of 15 or less that have their LOCATION attribute set to ABOVE, have the LOCATION attribute switched to BELOW and their PAGECLAS attribute set to 4KB and the buffer pool size is set to 1000 pages.
- Any cluster-sender channels that have been configured to use a transmission queue other than SYSTEM.CLUSTER.TRANSMIT.QUEUE fail to start with message CSQX295E. To allow these channels to start, you need to perform the following actions:
  - Change the default cluster transmission queue configuration of the queue manager, so that all cluster-sender channels default to use the transmission queue SYSTEM.CLUSTER.TRANSMIT.QUEUE. You can do this by changing the value of the DEFCLXQ queue manager attribute to SCTQ.
  - Identify any manually defined transmission queues that have a non blank cluster channel name attribute value, by using the DISPLAY QLOCAL(\*) WHERE(CLCHNAME NE ' ') command. Change the cluster channel-name attribute value of these queues to blank.
- No inbound channels will be allowed to start if any channel authentication records have been created with a host name specified in their **ADDRESS** attribute. Message CSQY344E is issued for each channel authentication rule that uses restricted function, if this condition occurs.
- Defining channel authentication (**CHLAUTH**) with the **CHKCLNT** attribute, requires the queue manager to be running in NEWFUNC mode.

If you need assistance to revert to an earlier version of IBM MQ, contact your IBM support center.

**Related information:**

Using CSQ6SYSP  
 DISPLAY SYSTEM

## PROPCTL channel options

Use **PROPCTL** channel attribute to control which message properties are included in a message that is sent from a Version 8.0 queue manager to a partner queue manager from an earlier version of IBM MQ.



Table 107. Channel message property attribute settings

| PROPCTL | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALL     | <p>Use this option if applications connected to the partner queue manager from an earlier version are able to process any properties placed in a message by a Version 8.0 application.</p> <p>All properties are sent to the partner queue manager, in addition to any name/value pairs placed in the MQRFH2.</p> <p>You must consider two application design issues:</p> <ol style="list-style-type: none"> <li>1. An application connected to the partner queue manager must be able to process messages containing MQRFH2 headers generated on a Version 8.0 queue manager.</li> <li>2. The application connected to the partner queue manager must process new message properties that are flagged with MQPD_SUPPORT_REQUIRED correctly.</li> </ol> <p>With the ALL channel option set, JMS applications can interoperate between IBM MQ Version 8.0 and an earlier version using the channel. New Version 8.0 applications using message properties can interoperate with applications from an earlier version, depending on how the earlier version application handles MQRFH2 headers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| COMPAT  | <p>Use this option to send message properties to applications connected to an earlier version partner queue manager in some cases, but not all. Message properties are only sent if two conditions are met:</p> <ol style="list-style-type: none"> <li>1. No property must be marked as requiring message property processing.</li> <li>2. At least one of the message properties must be in a “reserved” folder; see Note.</li> </ol> <p>With the COMPAT channel option set, JMS applications can interoperate between IBM MQ Version 8.0 and an earlier version using the channel.</p> <p>The channel is not available to every application using message properties, only to those applications that use the reserved folders. The rules concerning whether the message or the property is sent are:</p> <ol style="list-style-type: none"> <li>1. If the message has properties, but none of the properties are associated with a “reserved” folder, then no message properties are sent.</li> <li>2. If any message property has been created in a “reserved” property folder, all message properties associated with the message are sent. However: <ol style="list-style-type: none"> <li>a. If any of the message properties are marked as support being required, MQPD_SUPPORT_REQUIRED or MQPD_SUPPORT_REQUIRED_IF_LOCAL, the whole message is rejected. It is returned, discarded, or sent to the dead letter queue according to the value of its report options.</li> <li>b. If no message properties are marked as support being required, an individual property might not be sent. If any of the message property descriptor fields are set to non-default values the individual property is not sent. The message is still sent. An example of a non-default property descriptor field value is MQPD_USER_CONTEXT.</li> </ol> </li> </ol> <p><b>Note:</b> The “reserved” folders names start with mcd., jms., usr., or mqext.. These folders are created for applications that use the JMS interface. In Version 8.0 any name/value pairs that are placed in these folders are treated as message properties.</p> <p>Message properties are sent in an MQRFH2 header, in addition to any name/value pairs placed in an MQRFH2 header. Any name/value pairs placed in an MQRFH2 header are sent as long as the message is not rejected.</p> |
| NONE    | <p>Use this option to prevent any message properties being sent to applications connected to an earlier version partner queue manager. An MQRFH2 that contains name/value pairs and message properties is still sent, but only with the name/value pairs.</p> <p>With the NONE channel option set, a JMS message is sent as a JMSTextMessage or a JMSBytesMessage without any JMS message properties. If it is possible for an earlier version application to ignore all properties set in a Version 8.0 application, it can interoperate with it.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## PROPCTL queue options

Use the **PROPCTL** queue attribute to control how message properties are returned to an application that calls MQGET without setting any MQGMO message property options.

Table 108. Queue message property attribute settings

| PROPCTL | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALL     | <p>Use this option so that different applications reading a message from the same queue can process the message in different ways.</p> <ul style="list-style-type: none"><li>• An application, migrated unchanged from an earlier version, can continue to read the MQRFH2 directly. Properties are directly accessible in the MQRFH2 header.<br/>You must modify the application to handle any new properties, and new property attributes. It is possible that the application might be affected by changes in the layout and number of MQRFH2 headers. Some folder attributes might be removed, or that IBM MQ reports an error in the layout of the MQRFH2 header that it ignored in an earlier version.</li><li>• A new or changed application can use the message property MQI to query message properties, and read name/value pairs in MQRFH2 header directly.</li></ul> <p>All the properties in the message are returned to the application.</p> <ul style="list-style-type: none"><li>• If the application calls MQCRTMH to create a message handle, it must query the message properties using MQINQMP. Name/value pairs that are not message properties remain in the MQRFH2, which is stripped of any message properties.</li><li>• If the application does not create a message handle, all the message properties and name/value pairs remain in the MQRFH2.</li></ul> <p>ALL only has this affect if the receiving application has not set a MQGMO_PROPERTIES option, or has set it to MQGMO_PROPERTIES_AS_Q_DEF.</p> |

Table 108. Queue message property attribute settings (continued)

| PROPCTL       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>COMPAT</b> | <p>COMPAT is the default option. If <code>QMO_PROPERTIES_*</code> is not set, as in an unmodified application from an earlier version, COMPAT is assumed. By defaulting to the COMPAT option, an earlier version application that did not explicitly create an MQRFH2, works without change on Version 8.0.</p> <p>Use this option if you have written an earlier version application MQI application to read JMS messages.</p> <ul style="list-style-type: none"> <li>• The JMS properties, which are stored in an MQRFH2 header, are returned to the application in an MQRFH2 header in folders with names starting with <code>mcd.</code>, <code>jms.</code>, <code>usr.</code>, or <code>mqext.</code> .</li> <li>• If the message has JMS folders, and if a Version 8.0 application adds new property folders to the message, these properties are also returned in the MQRFH2. Consequently, you must modify the application to handle any new properties, and new property attributes. It is possible that an unmodified application might be affected by changes in the layout and number of MQRFH2 headers. It might find some folder attributes are removed, or that IBM MQ finds errors in the layout of the MQRFH2 header that it ignored in an earlier version.</li> </ul> <p><b>Note:</b> In this scenario, the behavior of the application is the same whether it is connected to an earlier version or Version 8.0 queue manager. If the channel <b>PROPCTL</b> attribute is set to COMPAT or ALL any new message properties are sent in the message to the earlier version partner queue manager.</p> <ul style="list-style-type: none"> <li>• If the message is not a JMS message, but contains other properties, those properties are not returned to the application in an MQRFH2 header. (The existence of specific property folders created by the IBM MQ classes for JMS indicates a JMS message. The property folders are <code>mcd.</code>, <code>jms.</code>, <code>usr.</code>, or <code>mqext.</code> )</li> <li>• The option also enables earlier version applications that explicitly create an MQRFH2 to work correctly, in many cases. For example, An MQI program that creates an MQRFH2 containing JMS message properties continues to work correctly. If a message is created without JMS message properties, but with some other MQRFH2 folders, the folders are returned to the application. Only if the folders are message property folders are those specific folders removed from the MQRFH2. Message property folders are identified by having the new folder attribute <code>content='properties'</code>, or are folders with names listed in Defined property folder name or Ungrouped property folder name.</li> <li>• If the application calls MQCRTMH to create a message handle, it must query the message properties using MQINQMP. Message properties are removed from the MQRFH2 headers. Name/value pairs that are not message properties remain in the MQRFH2.</li> <li>• If the application calls MQCRTMH to create a message handle, it can query all message properties, regardless of whether the message has JMS folders.</li> <li>• If the application does not create a message handle, all the message properties and name/value pairs remain in the MQRFH2.</li> </ul> <p>If a message contains new user property folders, you can infer that the message was created by a new or changed Version 8.0 application. If the receiving application is to process these new properties directly in an MQRFH2, you must modify the application to use the ALL option. With the default COMPAT option set, an unmodified application continues to process the rest of the MQRFH2, without the Version 8.0 properties.</p> <p>The intent of the PROPCTL interface is to support old applications reading MQRFH2 folders, and new and changed applications using the message property interface. Aim for new applications to use the message property interface for all user message properties, and to avoid reading and writing MQRFH2 headers directly.</p> <p>COMPAT only has this affect if the receiving application has not set a MQGMO_PROPERTIES option, or has set it to MQGMO_PROPERTIES_AS_Q_DEF.</p> |

Table 108. Queue message property attribute settings (continued)

| PROPCTL  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FORCE    | <p>The FORCE option places all message properties into MQRFH2 headers. All message properties and name/value pairs in the MQRFH2 headers remain in the message. Message properties are not removed from the MQRFH2, and made available through a message handle. The effect of choosing the FORCE option is to enable a newly migrated application to read message properties from MQRFH2 headers.</p> <p>Suppose you have modified an application to process Version 8.0 message properties, but have also retained its ability to work directly with MQRFH2 headers, as before. You can decide when to switch the application over to using message properties by initially setting the PROPCTL queue attribute to FORCE. Set the <b>PROPCTL</b> queue attribute to another value when you are ready to start using message properties. If the new function in the application does not behave as you expected, set the <b>PROPCTL</b> option back to FORCE.</p> <p>FORCE only has this affect if the receiving application has not set a MQGMO_PROPERTIES option, or has set it to MQGMO_PROPERTIES_AS_Q_DEF.</p> |
| NONE     | <p>Use this option so that an existing application can process a message, ignoring all message properties, and a new or changed application can query message properties.</p> <ul style="list-style-type: none"> <li>• If the application calls MQCRTMH to create a message handle, it must query the message properties using MQINQMP. Name/value pairs that are not message properties remain in the MQRFH2, which is stripped of any message properties.</li> <li>• If the application does not create a message handle, all the message properties are removed from the MQRFH2. Name/value pairs in the MQRFH2 headers remain in the message.</li> </ul> <p>NONE only has this affect if the receiving application has not set a MQGMO_PROPERTIES option, or has set it to MQGMO_PROPERTIES_AS_Q_DEF.</p>                                                                                                                                                                                                                                                                                                        |
| V6COMPAT | <p>Use this option to receive an MQRFH2 in the same format as it was sent. If the sending application, or the queue manager, creates additional message properties, these are returned in the message handle.</p> <p>This option has to be set on both the sending and receiving queues, and any intervening transmission queues. It overrides any other PROPCTL options set on queue definitions in the queue name resolution path.</p> <p>Use this option only in exceptional circumstances. For example, if you are migrating applications from an earlier version to Version 8.0 the option is valuable because it preserves the behavior of the earlier version. The option is likely to have an impact on message throughput. It is also more difficult to administer; you need to ensure the option is set on the sender, receiver, and intervening transmission queues.</p> <p>V6COMPAT only has this affect if the receiving application has not set a MQGMO_PROPERTIES option, or has set it to MQGMO_PROPERTIES_AS_Q_DEF.</p>                                                                             |

**Related information:**

PROPCTL

**MQGMO message property option settings**

Use MQGMO message property options to control how message properties are returned to an application.

Table 109. MQGMO message property option settings

| MQGMO Option               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQGMO_PROPERTIES_AS_Q_DEF  | <p>Applications prior to Version 7.0, and Version 8.0 applications that read from the same queue, and do not set <code>GMO_PROPERTIES_*</code>, receive the message properties differently. Applications prior to Version 7.0, and Version 8.0 applications that do not create a message handle, are controlled by the queue <b>PROPCTL</b> attribute. A Version 8.0 application can choose to receive message properties in the MQRFH2, or create a message handle and query the message properties. If the application creates a message handle, properties are removed from the MQRFH2.</p> <ul style="list-style-type: none"> <li>• An unmodified application prior to Version 7.0 does not set <code>GMO_PROPERTIES_*</code>. Any message properties it receives are in the MQRFH2 headers.</li> <li>• A new or changed Version 8.0 application that does not set <code>GMO_PROPERTIES_*</code> or sets it to <code>MQGMO_PROPERTIES_AS_Q_DEF</code> can choose to query message properties. It must set <code>MQCRTMH</code> to create a message handle and query message properties using the <code>MQINQMP</code> MQI call.</li> <li>• If a new or changed application does not create a message handle, it behaves like an application prior to Version 7.0. It must read any message properties it receives directly from the MQRFH2 headers.</li> <li>• If the queue attribute <b>PROPCTL</b> is set to <code>FORCE</code>, no properties are returned in the message handle. All properties are returned in MQRFH2 headers.</li> <li>• If the queue attribute <b>PROPCTL</b> is set to <code>NONE</code>, or <code>COMPAT</code>, a Version 8.0 application that creates a message handle, receives all message properties.</li> <li>• If the queue attribute <b>PROPCTL</b> is set to <code>V6COMPAT</code>, and also set to <code>V6COMPAT</code> on all the queues the message was placed on between the sender and the receiver, properties set by <code>MQSETMP</code> are returned in the message handle, and properties and name/value pairs created in an MQRFH2 are returned in the MQRFH2. The format of an MQRFH2 sent in Version 8.0 is the same as an application prior to Version 7.0, when sent by the same application.</li> </ul> |
| MQGMO_PROPERTIES_IN_HANDLE | <p>Force an application to use message properties. Use this option to detect if a modified application fails to create message handle. The application might be trying to read message properties directly from an MQRFH2, rather than call <code>MQINQMP</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MQGMO_NO_PROPERTIES        | <p>An application prior to Version 7.0 and a Version 8.0 application behave the same way, even if the Version 8.0 application creates a message handle.</p> <ul style="list-style-type: none"> <li>• All properties are removed. An unchanged application, prior to Version 7.0, connected to a Version 8.0 queue manager might behave differently to when it was connected to a partner queue manager from a version prior to Version 7.0. Queue manager generated properties, such as JMS properties, are removed.</li> <li>• Properties are removed even if a message handle is created. Name/value pairs in other MQRFH2 folders are available in the message data.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 109. MQGMO message property option settings (continued)

| MQGMO Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQGMO_PROPERTIES_FORCE_MQRFH2  | <p>Applications prior to Version 7.0 and Version 8.0 applications behave the same way. Properties are returned in the MQRFH2 headers, even if a message handle is created.</p> <ul style="list-style-type: none"> <li>• MQINQMP returns no message properties, even if a message handle is created. MQRC_PROPERTY_NOT_AVAILABLE is returned if a property is inquired upon.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| MQGMO_PROPERTIES_COMPATIBILITY | <p>An application prior to Version 7.0 connected to a Version 8.0 queue manager behaves the same as when it is connected to a queue manager from a version prior to Version 7.0. If the message is from a JMS client, the JMS properties are returned in the MQRFH2 headers. New or modified Version 8.0 applications, that create a message handle, behave differently.</p> <ul style="list-style-type: none"> <li>• All properties in any message property folders are returned if the message contains a mcd., jms., usr., or mqext. folder.</li> <li>• If the message contains property folders, but not a mcd., jms., usr., or mqext. folder, no message properties are returned in an MQRFH2.</li> <li>• If a new or modified Version 8.0 application creates a message handle, query message properties using the MQINQMP MQI call. All message properties are removed from the MQRFH2.</li> <li>• If a new or modified Version 8.0 application creates a message handle, all properties in the message can be queried. Even if the message does not contain a mcd., jms., usr., or mqext. folder, all message properties are queryable.</li> </ul> |

**Related information:**

PROPCTL

2471 (09A7) (RC2471): MQRC\_PROPERTY\_NOT\_AVAILABLE

---

## Configuring

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

Before you configure IBM MQ, read about the IBM MQ concepts in IBM MQ Technical overview. Read about how to plan your IBM MQ environment in Planning.

There are a number of different methods that you can use to create, configure, and administer your queue managers and their related resources in IBM MQ. These methods include command line interfaces, a graphical user interface, and an administration API. For more information about these interfaces, see *Administering IBM MQ*.

For instructions on how to create, start, stop, and delete a queue manager, see “Creating and managing queue managers on distributed platforms.”

For information about how to create the components required to connect your IBM MQ installations and applications together, see “Configuring distributed queuing” on page 802.

For instructions on how to connect your clients to an IBM MQ server by using different methods, see “Configuring connections between the server and client” on page 696.

For instructions on how to configure a queue manager cluster, see “Configuring a queue manager cluster” on page 902.

You can change the behavior of IBM MQ or a queue manager by changing configuration information. For more information, see “Changing IBM MQ and queue manager configuration information” on page 755. In general, you do not need to restart a queue manager for any configuration changes to take effect, except for when stated in this product documentation.

▶ **z/OS** For instructions on how to configure IBM MQ for z/OS, see “Configuring queue managers on z/OS” on page 1194.

### Related concepts:

▶ **z/OS** “Configuring queue managers on z/OS” on page 1194  
Use these instructions to configure queue managers on IBM MQ for z/OS.

### Related information:

IBM MQ technical overview

Administering local IBM MQ objects

Administering remote IBM MQ objects

▶ **IBM i** Administering IBM i

▶ **z/OS** Administering IBM MQ for z/OS

Planning

▶ **z/OS** Planning your IBM MQ environment on z/OS

---

## Creating and managing queue managers on distributed platforms

Before you can use messages and queues, you must create and start at least one queue manager and its associated objects.

## Creating a queue manager

A queue manager manages the resources associated with it, in particular the queues that it owns. It provides queuing services to applications for Message queuing Interface (MQI) calls and commands to create, modify, display, and delete IBM MQ objects.

To create a queue manager, you use the IBM MQ control command `crtmqm` (described in `crtmqm`). The `crtmqm` command automatically creates the required default objects and system objects (described in System default objects). Default objects form the basis of any object definitions that you make; system objects are required for queue manager operation. When you have created a queue manager and its objects, use the `strmqm` command to start the queue manager.

**Note:** IBM MQ does not support machine names that contain spaces. If you install IBM MQ on a computer with a machine name that contains spaces, you cannot create any queue managers.

Before you can create a queue manager, there are several points you must consider (especially in a production environment). Work through the following checklist:

### The installation associated with the queue manager

The `crtmqm` command automatically associates a queue manager with the installation from which the `crtmqm` command was issued. For commands that operate on a queue manager, you must issue the command from the installation associated with the queue manager. You can change the associated installation of a queue manager using the `setmqm` command. Note the Windows installer does not add the user that performs the install to the `mqm` group, for more details, see Authority to administer IBM MQ on UNIX, Linux and Windows systems.

### Naming conventions

Use uppercase names so that you can communicate with queue managers on all platforms. Remember that names are assigned exactly as you enter them. To avoid the inconvenience of lots of typing, do not use unnecessarily long names.

### Specify a unique queue manager name

When you create a queue manager, ensure that no other queue manager has the same name *anywhere* in your network. Queue manager names are not checked when the queue manager is created, and names that are not unique prevent you from creating channels for distributed queuing. Also, if you use the network for publish/subscribe messaging, subscriptions are associated with the queue manager name that created them. Therefore if queue managers in the cluster or hierarchy have the same name, it can result in publications not reaching them.

One way of ensuring uniqueness is to prefix each queue manager name with its own unique node name. For example, if a node is called ACCOUNTS, you can name your queue manager ACCOUNTS.SATURN.QUEUE.MANAGER, where SATURN identifies a particular queue manager and QUEUE.MANAGER is an extension you can give to all queue managers. Alternatively, you can omit this, but note that ACCOUNTS.SATURN and ACCOUNTS.SATURN.QUEUE.MANAGER are *different* queue manager names.

If you are using IBM MQ for communication with other enterprises, you can also include your own enterprise name as a prefix. This is not done in the examples, because it makes them more difficult to follow.

**Note:** Queue manager names in control commands are case-sensitive. This means that you are allowed to create two queue managers with the names `jupiter.queue.manager` and `JUPITER.queue.manager`. However, it is better to avoid such complications.

### Limit the number of queue managers

You can create as many queue managers as resources allow. However, because each queue manager requires its own resources, it is generally better to have one queue manager with 100 queues on a node than to have ten queue managers with ten queues each.



In production systems, many processors can be exploited with a single queue manager, but larger server machines might run more effectively with multiple queue managers.

### Specify a default queue manager

Each node should have a default queue manager, though it is possible to configure IBM MQ on a node without one. The default queue manager is the queue manager that applications connect to if they do not specify a queue manager name in an **MQCONN** call. It is also the queue manager that processes MQSC commands when you invoke the **runmqsc** command without specifying a queue manager name.

Specifying a queue manager as the default *replaces* any existing default queue manager specification for the node.

Changing the default queue manager can affect other users or applications. The change has no effect on currently-connected applications, because they can use the handle from their original connect call in any further MQI calls. This handle ensures that the calls are directed to the same queue manager. Any applications connecting *after* you have changed the default queue manager connect to the new default queue manager. This might be what you intend, but you should take this into account before you change the default.

Creating a default queue manager is described in “Creating a default queue manager” on page 690.

### Specify a dead-letter queue

The dead-letter queue is a local queue where messages are put if they cannot be routed to their intended destination.

It is important to have a dead-letter queue on each queue manager in your network. If you do not define one, errors in application programs might cause channels to be closed, and replies to administration commands might not be received.

For example, if an application tries to put a message on a queue on another queue manager, but gives the wrong queue name, the channel is stopped and the message remains on the transmission queue. Other applications cannot then use this channel for their messages.

The channels are not affected if the queue managers have dead-letter queues. The undelivered message is put on the dead-letter queue at the receiving end, leaving the channel and its transmission queue available.

When you create a queue manager, use the **-u** flag to specify the name of the dead-letter queue. You can also use an MQSC command to alter the attributes of a queue manager that you have already defined to specify the dead-letter queue to be used. See *Working with queue managers* for an example of the MQSC command **ALTER**.

### Specify a default transmission queue

A transmission queue is a local queue on which messages in transit to a remote queue manager are queued before transmission. The default transmission queue is the queue that is used when no transmission queue is explicitly defined. Each queue manager can be assigned a default transmission queue.

When you create a queue manager, use the **-d** flag to specify the name of the default transmission queue. This does not actually create the queue; you have to do this explicitly later on. See *Working with local queues* for more information.

### Specify the logging parameters you require

You can specify logging parameters on the **crtmqm** command, including the type of logging, and the path and size of the log files.

In a development environment, the default logging parameters should be adequate. However, you can change the defaults if, for example:

- You have a low-end system configuration that cannot support large logs.
- You anticipate a large number of long messages being on your queues at the same time.

- You anticipate a lot of persistent messages passing through the queue manager.

Once you have set the logging parameters, some of them can only be changed by deleting the queue manager and recreating it with the same name but with different logging parameters.

For more information about logging parameters, see “Availability, recovery and restart” on page 1037.

UNIX

#### For IBM MQ for UNIX systems only

You can create the queue manager directory `/var/mqm/qmgrs/<qmgr>`, even on a separate local file system, before you use the `crtmqm` command. When you use `crtmqm`, if the `/var/mqm/qmgrs/<qmgr>` directory exists, is empty, and is owned by `mqm`, it is used for the queue manager data. If the directory is not owned by `mqm`, the creation fails with a First Failure Support Technology (FFST) message. If the directory is not empty, a new directory is created.

#### Related concepts:

z/OS

“Configuring queue managers on z/OS” on page 1194

Use these instructions to configure queue managers on IBM MQ for z/OS.

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Backing up configuration files after creating a queue manager” on page 692

IBM MQ configuration information is stored in configuration files on Windows, UNIX and Linux systems.

“Starting a queue manager” on page 693

When you create a queue manager, you must start it to enable it to process commands or MQI calls.

“Stopping a queue manager” on page 694

There are three ways to stop a queue manager: a quiesced shutdown, and immediate shutdown, and a preemptive shutdown.

“Restarting a queue manager” on page 695

You can use the `strmqm` command to restart a queue manager, or on IBM MQ for Windows and IBM MQ for Linux x86-64 systems, restart a queue manager from IBM MQ Explorer.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

#### Related tasks:

“Making an existing queue manager the default” on page 692

You can make an existing queue manager the default queue manager. The way you do this depends on the platform you are using.

“Deleting a queue manager” on page 695

You can delete a queue manager using the `dltmqm` command or by using the IBM MQ Explorer.

#### Related information:

Creating a queue manager called QM1

System and default objects

## Creating a default queue manager

distributed

The default queue manager is the queue manager that applications connect to if they do not specify a queue manager name in an MQCONN call. It is also the queue manager that processes MQSC commands when you invoke the `runmqsc` command without specifying a queue manager name. To create a queue manager, you use the IBM MQ control command `crtmqm`.

## Before you begin

Before creating a default queue manager, read through the considerations described in “Creating and managing queue managers on distributed platforms” on page 687.

**UNIX** When you use **crtmqm** to create a queue manager on UNIX, if the `/var/mqm/qmgrs/<qmgr>` directory already exists, is owned by `mqm`, and is empty, it is used for the queue manager data. If the directory is not owned by `mqm`, the creation of the queue manager fails with a First Failure Support Technology (FFST) message. If the directory is not empty, a new directory is created for the queue manager data.

This consideration applies even when the `/var/mqm/qmgrs/<qmgr>` directory already exists on a separate local file system.

## About this task

When you create a queue manager by using the **crtmqm** command, the command automatically creates the required default objects and system objects. Default objects form the basis of any object definitions that you make and system objects are required for queue manager operation.

By including the relevant parameters in the command, you can also define, for example, the name of the default transmission queue to be used by the queue manager, and the name of the dead letter queue.

**Windows** On Windows, you can use the **sax** option of the **crtmqm** command to start multiple instances of the queue manager.

For more information about the **crtmqm** command and its syntax, see **crtmqm**.

## Procedure

To create a default queue manager, use the **crtmqm** command with the **-q** flag. The following example of the **crtmqm** command creates a default queue manager called `SATURN.QUEUE.MANAGER`:

```
crtmqm -q -d MY.DEFAULT.XMIT.QUEUE -u SYSTEM.DEAD.LETTER.QUEUE SATURN.QUEUE.MANAGER
```

where:

**-q** Indicates that this queue manager is the default queue manager.

**-d MY.DEFAULT.XMIT.QUEUE**

Is the name of the default transmission queue to be used by this queue manager.

**Note:** IBM MQ does not create a default transmission queue for you; you have to define it yourself.

**-u SYSTEM.DEAD.LETTER.QUEUE**

Is the name of the default dead-letter queue created by IBM MQ on installation.

**SATURN.QUEUE.MANAGER**

Is the name of this queue manager. This must be the last parameter specified on the **crtmqm** command.

## What to do next

When you have created a queue manager and its objects, use the **strmqm** command to start the queue manager.

### Related concepts:

“Backing up configuration files after creating a queue manager”

IBM MQ configuration information is stored in configuration files on Windows, UNIX and Linux systems.

### Related information:

Working with queue managers

Working with local queues

System and default objects

## Making an existing queue manager the default

You can make an existing queue manager the default queue manager. The way you do this depends on the platform you are using.

### IBM MQ for Windows and IBM MQ for Linux (x86 and x86-64 platforms) systems

#### About this task

Use the following instructions to make an existing queue manager the default queue manager on IBM MQ for Windows and IBM MQ for Linux (x86 and x86-64 platforms) systems:

#### Procedure

1. Open the IBM MQ Explorer.
2. Right-click **IBM WebSphere MQ** , then select **Properties...** . The Properties for IBM MQ panel is displayed.
3. Type the name of the default queue manager into the Default queue manager name field.
4. Click OK.

### UNIX and Linux systems

#### About this task

When you create a default queue manager, its name is inserted in the Name attribute of the `DefaultQueueManager` stanza in the IBM MQ configuration file (`mqs.ini`). The stanza and its contents are automatically created if they do not exist.

#### Procedure

- To make an existing queue manager the default, change the queue manager name on the Name attribute to the name of the new default queue manager. You can do this manually, using a text editor.
- If you do not have a default queue manager on the node, and you want to make an existing queue manager the default, create the `DefaultQueueManager` stanza with the required name yourself.
- If you accidentally make another queue manager the default and want to revert to the original default queue manager, edit the `DefaultQueueManager` stanza in `mqs.ini`, replacing the unwanted default queue manager with that of the one you want.

#### What to do next

See “Changing IBM MQ and queue manager configuration information” on page 755 for information about configuration files.

## Backing up configuration files after creating a queue manager

IBM MQ configuration information is stored in configuration files on Windows, UNIX and Linux systems.

On Windows and Linux (x86 and x86-64) systems use MQ Explorer to make changes to the configuration files.

On Windows systems you can also use the **amqmdain** command to make changes to the configuration files. See, amqmdain

There are two types of configuration file:

- When you install the product, the IBM MQ configuration file (mqs.ini) is created. It contains a list of queue managers that is updated each time you create or delete a queue manager. There is one mqs.ini file per node.
- When you create a new queue manager, a new queue manager configuration file (qm.ini) is automatically created. This contains configuration parameters for the queue manager.

**V8.0.0.4** If you have installed the AMQP service, then there is an additional configuration file that you must back up:

- amqp\_win.properties (Windows)
- amqp\_unix.properties (UNIX/Linux)

After creating a queue manager, back up your configuration files. Then, if you create another queue manager that causes you problems, you can reinstate the backups when you have removed the source of the problem. As a general rule, back up your configuration files each time you create a new queue manager.

For more information about configuration files, see “Changing IBM MQ and queue manager configuration information” on page 755.

## Starting a queue manager

When you create a queue manager, you must start it to enable it to process commands or MQI calls.

To start a queue manager, use the **strmqm** command.

**Note:** You must use the **strmqm** command from the installation associated with the queue manager that you are working with. You can find out which installation a queue manager is associated with using the `dspmqs -o` installation command.

For example, to start a queue manager QMB enter the following command:

```
strmqm QMB
```

On IBM MQ for Windows and IBM MQ for Linux (x86 and x86-64 platforms) systems, you can start a queue manager as follows:

1. Open the IBM MQ Explorer.
2. Select the queue manager from the Navigator View.
3. Click **Start** . The queue manager starts.

If the queue manager start-up takes more than a few seconds IBM MQ issues information messages intermittently detailing the start-up progress.

The **strmqm** command does not return control until the queue manager has started and is ready to accept connection requests.

## Starting a queue manager automatically

In IBM MQ for Windows you can start a queue manager automatically when the system starts using the IBM MQ Explorer. For more information, see Administration using the MQ Explorer .

## Stopping a queue manager

There are three ways to stop a queue manager: a quiesced shutdown, and immediate shutdown, and a preemptive shutdown.

Use the **endmqm** command to stop a queue manager.

**Note:** You must use the **endmqm** command from the installation associated with the queue manager that you are working with. You can find out which installation a queue manager is associated with using the `dspmqr -o` installation command.

For example, to stop a queue manager called QMB, enter the following command:

```
endmqm QMB
```

On IBM MQ for Windows and IBM MQ for Linux (x86 and x86-64 platforms) systems, you can stop a queue manager as follows:

1. Open the IBM MQ Explorer.
2. Select the queue manager from the Navigator View.
3. Click **Stop...** . The End Queue Manager panel is displayed.
4. Select Controlled, or Immediate.
5. Click **OK** . The queue manager stops.

### Quiesced shutdown

By default, the **endmqm** command performs a quiesced shutdown of the specified queue manager. This might take a while to complete. A quiesced shutdown waits until all connected applications have disconnected.

Use this type of shutdown to notify applications to stop. If you issue:

```
endmqm -c QMB
```

you are not told when all applications have stopped. (An `endmqm -c QMB` command is equivalent to an `endmqm QMB` command.)

However, if you issue:

```
endmqm -w QMB
```

the command waits until all applications have stopped and the queue manager has ended.

### Immediate shutdown

For an immediate shutdown any current MQI calls are allowed to complete, but any new calls fail. This type of shutdown does not wait for applications to disconnect from the queue manager.

For an immediate shutdown, type:

```
endmqm -i QMB
```

### Preemptive shutdown

**Note:** Do not use this method unless all other attempts to stop the queue manager using the **endmqm** command have failed. This method can have unpredictable consequences for connected applications.

If an immediate shutdown does not work, you must resort to a *preemptive* shutdown, specifying the `-p` flag. For example:

```
endmqm -p QMB
```

This stops the queue manager immediately. If this method still does not work, see [Stopping a queue manager manually](#) for an alternative solution.

For a detailed description of the **endmqm** command and its options, see [endmqm](#).

## If you have problems shutting down a queue manager

Problems in shutting down a queue manager are often caused by applications. For example, when applications:

- Do not check MQI return codes properly
- Do not request notification of a quiesce
- Terminate without disconnecting from the queue manager (by issuing an **MQDISC** call)

If a problem occurs when you stop the queue manager, you can break out of the **endmqm** command using Ctrl-C. You can then issue another **endmqm** command, but this time with a flag that specifies the type of shutdown that you require.

## Restarting a queue manager

You can use the **strmqm** command to restart a queue manager, or on IBM MQ for Windows and IBM MQ for Linux x86-64 systems, restart a queue manager from IBM MQ Explorer.

To restart a queue manager, type:

```
strmqm saturn.queue.manager
```

On IBM MQ for Windows and IBM MQ for Linux x86-64 systems, you can restart a queue manager in the same way as starting it, as follows:

1. Open the IBM MQ Explorer.
2. Select the queue manager from the Navigator View.
3. Click **Start** . The queue manager restarts.

If the queue manager restart takes more than a few seconds IBM MQ issues information messages intermittently detailing the start-up progress.

## Deleting a queue manager

You can delete a queue manager using the **dltmqm** command or by using the IBM MQ Explorer.

### Before you begin

Stop the queue manager.

### Procedure

Issue the following command: **dltmqm QMB**

**Note:** You must use the **dltmqm** command from the installation associated with the queue manager that you are working with. You can find out which installation a queue manager is associated with using the **dspmqs -o** installation command.

## Steps for deleting a queue manager

### About this task

On IBM MQ for Windows and IBM MQ for Linux (x86 and x86-64 platforms) systems, you can delete a queue manager as follows:

### Procedure

1. Open the IBM MQ Explorer.
2. In the Navigator view, select the queue manager.
3. If the queue manager is not stopped, stop it.
  - a. Right-click the queue manager.
  - b. Click **Stop**.
4. Right-click the queue manager.
5. Click **Delete**.

### Results

The queue manager is deleted.

### Attention:

- Deleting a queue manager is a drastic step, because you also delete all resources associated with the queue manager, including all queues and their messages and all object definitions. If you use the **dltmqm** command, there is no displayed prompt that allows you to change your mind; when you press the Enter key all the associated resources are lost.
- In IBM MQ for Windows, deleting a queue manager also removes the queue manager from the automatic startup list (described in “Starting a queue manager” on page 693 ). When the command has completed, a IBM MQ queue manager ending message is displayed; you are not told that the queue manager has been deleted.
- Deleting a cluster queue manager does not remove it from the cluster. See the note in the description of **dltmqm** for more information.

For a description of the **dltmqm** command and its options, see `dltmqm`. Ensure that only trusted administrators have the authority to use this command. (For information about security, see Setting up security on Windows, UNIX and Linux systems.)

---

## Configuring connections between the server and client

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

In IBM MQ, the logical communication links between objects are called *channels*. The channels used to connect IBM MQ MQI clients to servers are called MQI channels. You set up channel definitions at each end of your link so that your IBM MQ application on the IBM MQ MQI client can communicate with the queue manager on the server. For a detailed description of how to do this, see User defined channels.

Before you define your MQI channels, you must:

1. Decide on the form of communication that you are going to use. See “Which communication type to use” on page 697.
2. Define the connection at each end of the channel:

To define the connection, you must:

  - Configure the connection.



- Record the values of the parameters that you need for the channel definitions.
- Enable the server to detect incoming network requests from your IBM MQ MQI client, by starting a *listener*.

## Which communication type to use

Different platforms support different transmission protocols. Your choice of transmission protocol depends on your combination of IBM MQ MQI client and server platforms.





There are up to four types of transmission protocol for MQI channels depending on your client and server platforms:

- LU 6.2
- NetBIOS
- SPX
- TCP/IP

When you define your MQI channels, each channel definition must specify a transmission protocol (transport type) attribute. A server is not restricted to one protocol, so different channel definitions can specify different protocols. For IBM MQ MQI clients, it might be useful to have alternative MQI channels using different transmission protocols.

Your choice of transmission protocol might be restricted by your particular combination of IBM MQ MQI client and server platforms. The possible combinations are shown in the following table.





*Table 110. Transmission protocols - combination of IBM MQ MQI client and server platforms*

| Transmission protocol | IBM MQ MQI client                                                                                                 | IBM MQ server                                                                                                                                                                                                              |
|-----------------------|-------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TCP/IP                |  IBM i UNIX systems<br>Windows |  IBM i UNIX systems<br>Windows z/OS                                                                                                    |
| LU 6.2                | UNIX systems <sup>1</sup><br>Windows                                                                              |  IBM i UNIX systems <sup>1</sup><br>Windows  z/OS |
| NetBIOS               | Windows                                                                                                           | Windows                                                                                                                                                                                                                    |
| SPX                   | Windows                                                                                                           | Windows                                                                                                                                                                                                                    |

### Note:

1. Except Linux for Power Systems™

For more about setting up different types of connections, see the following links:

- “TCP/IP connection limits” on page 700
- “Defining a TCP connection on Windows” on page 866
- “Defining a TCP connection on UNIX and Linux” on page 873
-  “Defining a TCP connection on IBM i” on page 894
-  “Defining a TCP connection on z/OS” on page 1284
- “Defining an LU 6.2 connection on Windows” on page 868
- “Defining an LU 6.2 connection on UNIX and Linux” on page 877
-  “Defining an LU 6.2 connection on IBM i” on page 896
-  “Defining an LU6.2 connection for z/OS using APPC/MVS” on page 1287

- “Defining a NetBIOS connection on Windows” on page 870

**Related concepts:**

“Configuring an extended transactional client” on page 701

This collection of topics describes how to configure the extended transactional function for each category of transaction manager.

“Defining MQI channels” on page 712

To create a new channel, you have to create **two** channel definitions, one for each end of the connection, using the same channel name and compatible channel types. In this case, the channel types are *server-connection* and *client-connection*.

“Creating server-connection and client-connection definitions on different platforms” on page 713


You can create each channel definition on the computer to which it applies. There are restrictions on how you can create channel definitions on a client computer.

“Creating server-connection and client-connection definitions on the server” on page 716

You can create both definitions on the server, then make the client-connection definition available to the client.

“Channel-exit programs for MQI channels” on page 721

Three types of channel exit are available to the IBM MQ MQI client environment on UNIX, Linux and Windows systems.

 “Connecting a client to a queue-sharing group” on page 726

You can connect a client to a queue-sharing group by creating an MQI channel between a client and a queue manager on a server that is a member of a queue-sharing group.

“Configuring a client using a configuration file” on page 727

Configure your clients by using attributes in a text file. These attributes can be overridden by environment variables or in other platform-specific ways.

**Related information:**

Connecting IBM MQ MQI client applications to queue managers

DISPLAY CHLAUTH

SET CHLAUTH

## Which communication type to use

Different platforms support different communication protocols. Your choice of transmission protocol depends on your combination of IBM MQ MQI client and server platforms.

There are four types of communication for MQI channels on different platforms:

- LU 6.2
- NetBIOS
- SPX
- TCP/IP






When you define your MQI channels, each channel definition must specify a transmission protocol (transport type) attribute. A server is not restricted to one protocol, so different channel definitions can specify different protocols. For IBM MQ MQI clients, it might be useful to have alternative MQI channels using different transmission protocols.

Your choice of transmission protocol also depends on your particular combination of IBM MQ client and server platforms. The possible combinations are shown in the following table.

**Note:**

1. Except Linux ( POWER platform)

Table 111. Transmission protocols - combination of IBM MQ client and server platforms

| Transmission protocol | IBM MQ MQI client                                                                                               | IBM MQ server                                                                                                                                                                                                          |
|-----------------------|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TCP/IP                |  IBM i UNIX systems<br>Windows |  IBM i UNIX systems<br>Windows  z/OS              |
| LU 6.2                | UNIX systems <sup>1</sup><br>Windows                                                                            |  IBM i UNIX systems <sup>1</sup><br>Windows  z/OS |
| NetBIOS               | Windows                                                                                                         | Windows                                                                                                                                                                                                                |
| SPX                   | Windows                                                                                                         | Windows                                                                                                                                                                                                                |


**Related concepts:**

“Defining a TCP connection on Windows” on page 866

Define a TCP connection by configuring a channel at the sending end to specify the address of the target, and by running a listener program at the receiving end.

“Defining a TCP connection on UNIX and Linux” on page 873

The channel definition at the sending end specifies the address of the target. The listener or inet daemon is configured for the connection at the receiving end.

 “Defining a TCP connection on IBM i” on page 894

You can define a TCP connection within the channel definition using the Connection Name field.

 “Defining a TCP connection on z/OS” on page 1284


To define a TCP connection, there are a number of settings to configure.

“Defining an LU 6.2 connection on Windows” on page 868

SNA must be configured so that an LU 6.2 conversation can be established between the two machines.

“Defining an LU 6.2 connection on UNIX and Linux” on page 877

SNA must be configured so that an LU 6.2 conversation can be established between the two machines.

 “Defining an LU 6.2 connection on IBM i” on page 896

Define the LU 6.2 communications details by using a mode name, TP name, and connection name of a fully qualified LU 6.2 connection.

“Defining a NetBIOS connection on Windows” on page 870


IBM MQ uses three types of NetBIOS resource when establishing a NetBIOS connection to another IBM MQ product: sessions, commands, and names. Each of these resources has a limit, which is established either by default or by choice during the installation of NetBIOS.

**Related reference:**

“TCP/IP connection limits” on page 700

The number of outstanding connection requests that can be queued at a single TCP/IP port depends on the platform. An error occurs if the limit is reached.

**Related information:**

 “Defining an LU6.2 connection for z/OS using APPC/MVS” on page 1287

To define an LU6.2 connection there are a number of settings to configure.

**Defining a TCP/IP connection**

Specifying a transport type of TCP on the channel definition at the client end. Start a listener program on the server.

Specify a TCP/IP connection at the client by specifying a transport type of TCP on the channel definition.

Receiving channel programs are started in response to a startup request from the sending channel. To do this, a listener program has to be started to detect incoming network requests and start the associated channel. The procedure for starting a listener program depends on the server platform.

See the related topics for your client and server platforms.

### TCP/IP connection limits

The number of outstanding connection requests that can be queued at a single TCP/IP port depends on the platform. An error occurs if the limit is reached.

This connection limit is not the same as the maximum number of clients you can attach to an IBM MQ server. You can connect more clients to a server, up to the level determined by the server system resources. The backlog values for connection requests are shown in the following table:

*Table 112. Maximum outstanding connection requests queued at a TCP/IP port*

| Server platform     | Maximum connection requests |
|---------------------|-----------------------------|
| AIX                 | 100                         |
| HP-UX               | 20                          |
| Linux               | 100                         |
| IBM                 | 255                         |
| Solaris             | 100                         |
| Windows Server      | 100                         |
| Windows Workstation | 100                         |
| z/OS                | 255                         |

If the connection limit is reached, the client receives a return code of MQRC\_HOST\_NOT\_AVAILABLE from the **MQCONN** call, and an AMQ9202 error in the client error log ( /var/mqm/errors/AMQERR0n.LOG on UNIX and Linux systems or amqerr0n.log in the errors subdirectory of the IBM MQ client installation on Windows ). If the client retries the **MQCONN** request, it might be successful.

To increase the number of connection requests you can make, and avoid error messages being generated by this limitation, you can have multiple listeners each listening on a different port, or have more than one queue manager.

### Defining a NetBIOS connection

NetBIOS connections apply only to Windows systems.

A NetBIOS connection applies only to a client and server running Windows. See Defining a NetBIOS connection.

## Configuring an extended transactional client

This collection of topics describes how to configure the extended transactional function for each category of transaction manager.

For each platform, the extended transactional client provides support for the following external transaction managers:

### XA-compliant transaction managers

The extended transactional client provides the XA resource manager interface to support XA-compliant transaction managers such as CICS and Tuxedo.

### Microsoft Transaction Server ( Windows systems only)

On Windows systems only, the XA resource manager interface also supports Microsoft Transaction Server (MTS). The IBM MQ MTS support supplied with the extended transactional client provides the bridge between MTS and the XA resource manager interface.

### WebSphere Application Server

Earlier versions of IBM MQ supported WebSphere Application Server Version 4 or Version 5, and required you to carry out certain configuration tasks to use the extended transactional client. WebSphere Application Server Version 6 and later includes an IBM MQ messaging provider, so you do not need to use the extended transactional client.

### Related concepts:

“Configuring XA-compliant transaction managers”

First configure the IBM MQ base client, then configure the extended transactional function using the information in these topics.

“Microsoft Transaction Server” on page 711

No additional configuration is required before you can use MTS as a transaction manager. However, there are some points to note.

## Configuring XA-compliant transaction managers

First configure the IBM MQ base client, then configure the extended transactional function using the information in these topics.

**Note:** This section assumes that you have a basic understanding of the XA interface as published by The Open Group in *Distributed Transaction Processing: The XA Specification*.

To configure an extended transactional client, you must first configure the IBM MQ base client as described in Installing an IBM MQ client. Using the information in this section, you can then configure the extended transactional function for an XA-compliant transaction manager such as CICS and Tuxedo.

A transaction manager communicates with a queue manager as a resource manager using the same MQI channel as that used by the client application that is connected to the queue manager. When the transaction manager issues a resource manager (xa\_) function call, the MQI channel is used to forward the call to the queue manager, and to receive the output back from the queue manager.

Either the transaction manager can start the MQI channel by issuing an xa\_open call to open the queue manager as a resource manager, or the client application can start the MQI channel by issuing an MQCONN or MQCONNX call.

- If the transaction manager starts the MQI channel, and the client application later calls MQCONN or MQCONNX on the same thread, the MQCONN or MQCONNX call completes successfully and a connection handle is returned to the application. The application does not receive a MQCC\_WARNING completion code with an MQRC\_ALREADY\_CONNECTED reason code.
- If the client application starts the MQI channel, and the transaction manager later calls xa\_open on the same thread, the xa\_open call is forwarded to the queue manager using the MQI channel.

In a recovery situation following a failure, when no client applications are running, the transaction manager can use a dedicated MQI channel to recover any incomplete units of work in which the queue manager was participating at the time of the failure.

Note the following conditions when using an extended transactional client with an XA-compliant transaction manager:

- Within a single thread, a client application can be connected to only one queue manager at a time. This restriction applies only when using an extended transactional client; a client application that is using an IBM MQ base client can be connected to more than one queue manager concurrently within a single thread.
- Each thread of a client application can connect to a different queue manager.
- A client application cannot use shared connection handles.

To configure the extended transactional function, you must provide the following information to the transaction manager for each queue manager that acts as a resource manager:

- An `xa_open` string
- A pointer to an XA switch structure

When the transaction manager calls `xa_open` to open the queue manager as a resource manager, it passes the `xa_open` string to the extended transactional client as the argument, `xa_info`, on the call. The extended transactional client uses the information in the `xa_open` string in the following ways:

- To start an MQI channel to the server queue manager, if the client application has not already started one
- To check that the queue manager that the transaction manager opens as a resource manager is the same as the queue manager to which the client application connects
- To locate the transaction manager's `ax_reg` and `ax_unreg` functions, if the queue manager uses dynamic registration

For the format of an `xa_open` string, and for more details about how the information in the `xa_open` string is used by an extended transactional client, see “The format of an `xa_open` string” on page 704.

An XA switch structure enables the transaction manager to locate the `xa_` functions provided by the extended transactional client, and specifies whether the queue manager uses dynamic registration. For information about the XA switch structures supplied with an extended transactional client, see “The XA switch structures” on page 708.

For information about how to configure the extended transactional function for a particular transaction manager, and for any other information about using the transaction manager with an extended transactional client, see the following sections:

- “Configuring an extended transactional client for CICS” on page 710
- “Configuring an extended transactional client for Tuxedo” on page 711

**Related concepts:**

“The CHANNEL, TRPTYPE, CONNAME, and QMNAME parameters of the xa\_open string” on page 706  
Use this information to understand how the extended transactional client uses these parameters to determine the queue manager to connect to.

“Additional error processing for xa\_open” on page 707

The xa\_open call fails in certain circumstances.

**Related tasks:**

“Using the extended transactional client with SSL channels” on page 709

You cannot set up an SSL channel using the xa\_open string. Follow these instructions to use the client channel definition table (ccdt).

**Related reference:**

“The TPM and AXLIB parameters” on page 707

An extended transactional client uses the TPM and AXLIB parameters to locate the transaction manager's ax\_reg and ax\_unreg functions. These functions are used only if the queue manager uses dynamic registration.

“Recovery following a failure in extended transactional processing” on page 708

Following a failure, a transaction manager must be able to recover any incomplete units of work. To do this, the transaction manager must be able to open as a resource manager any queue manager that was participating in an incomplete unit of work at the time of the failure.

**IBM WebSphere MQ for z/OS considerations for extended transactional client connections:**

Some XA transaction managers use sequences of transaction coordination calls which are incompatible with the features normally available to clients connecting to IBM WebSphere MQ for z/OS.

Where an incompatible sequence is detected, IBM WebSphere MQ for z/OS might issue an abend for the connection and return an error response to the client.

For example, xa\_prepare receives abend 5C6-00D4007D, with return code -3 (XAER\_RMERR) returned to the client.

For transaction managers which encounter this situation, take the following actions to allow the transaction manager to interact with IBM WebSphere MQ for z/OS:

- Apply the fix for APAR PI49236.
- Enable the change provided by PI49236 for the server-connection channel used by the transaction manager.

You enable the change by specifying the keyword CSQSERVICE1 (in upper case) anywhere in the description field of the SVRCONN channel.

Note that channels with the CSQSERVICE1 keyword have the following restrictions:

- GROUP unit of recovery disposition is not permitted. Only QMGR unit of recovery disposition is allowed.

An xa\_open call specifying the queue sharing group name in the **xa\_info** parameter fails with *xaer\_inval*.

- The MQGMO\_LOCK and MQGMO\_UNLOCK options are not permitted. An MQGET call with MQGMO\_LOCK or MQGMO\_UNLOCK fails with MQRC\_ENVIRONMENT\_ERROR.

## Related concepts:

“Configuring XA-compliant transaction managers” on page 701

First configure the IBM MQ base client, then configure the extended transactional function using the information in these topics.

## The format of an xa\_open string:

An xa\_open string contains pairs of defined parameter names and values.

An xa\_open string has the following format:

```
parm_name1 = parm_value1, parm_name2 = parm_value2, ...
```

where *parm\_name* is the name of a parameter and *parm\_value* is the value of a parameter. The names of the parameters are not case-sensitive but, unless stated otherwise, the values of the parameters are case-sensitive. You can specify the parameters in any order.

The names, meanings, and valid values of the parameters are as follows:

### Name Meaning and valid values

#### CHANNEL

The name of an MQI channel.

This is an optional parameter. If this parameter is supplied, the CONNAME parameter must also be supplied.

#### TRPTYPE

The communications protocol for the MQI channel. The following are valid values:

**LU62** SNA LU 6.2

#### NETBIOS

NetBIOS

**SPX** IPX/SPX

**TCP** TCP/IP

This is an optional parameter. If it is omitted, the default value of TCP is assumed. The values of the parameter are not case-sensitive.

#### CONNAME

The network address of the queue manager at the server end of the MQI channel. The valid values of this parameter depend on the value of the TRPTYPE parameter:

**LU62** A symbolic destination name, which identifies a CPI-C side information entry.

The network qualified name of a partner LU is not a valid value, nor is a partner LU alias. This is because there are no additional parameters to specify a transaction program (TP) name and a mode name.

#### NETBIOS

A NetBIOS name.

**SPX** A 4 byte network address, a 6 byte node address, and an optional 2 byte socket number. These values must be specified in hexadecimal notation. A period must separate the network and node addresses, and the socket number, if supplied, must be enclosed in parentheses. For example:

```
0a0b0c0d.804abcde23a1(5e86)
```

If the socket number is omitted, the default value of 5e86 is assumed.



**TCP** A host name or an IP address, optionally followed by a port number in parentheses. If the port number is omitted, the default value of 1414 is assumed.

This is an optional parameter. If this parameter is supplied, the CHANNEL parameter must also be supplied.

### QMNAME

The name of the queue manager at the server end of the MQI channel. The name cannot be blank or a single asterisk (\*), nor can the name start with an asterisk. This means that the parameter must identify a specific queue manager by name.

This is a mandatory parameter.

When a client application is connected to a specific queue manager any transaction recovery must be processed by the same queue manager.

If the application is connecting to a z/OS queue manager then the application can specify either the name of a specific queue manager or the name of a queue-sharing group (QSG). By using the queue manager name or QSG name, the application controls whether it partakes in a transaction with a QMGR unit of recovery disposition or a GROUP unit of recovery disposition. The GROUP unit of recovery disposition enables the recovery of the transaction to be processed on any member of the QSG. To use GROUP units of recovery the **GROUPUR** queue manager attribute must be enabled.

 For further information about using GROUP unit of recovery, see Unit of recovery disposition in a queue-sharing group.

**TPM** The transaction manager being used. The valid values are CICS and TUXEDO.

An extended transactional client uses this parameter and the AXLIB parameter for the same purpose. For more information these parameters, see The TPM and AXLIB parameters.

This is an optional parameter. The values of the parameter are not case-sensitive.

### AXLIB

The name of the library that contains the transaction manager's ax\_reg and ax\_unreg functions.

This is an optional parameter.

**UID** The user ID that is provided to the queue manager for authentication. If this parameter is supplied, the **PWD** parameter must also be supplied. If the user ID and password supplied are authenticated, the user ID is used for identification of the transaction manager's connection. The user ID and password populate the MQCSP object on the MQCONN call.

The **UID** and **PWD** parameters are valid for both client and server bindings.

**PWD** The password that is provided to the queue manager for authentication. If this parameter is supplied, the **UID** parameter must also be supplied.

**Warning:** In some cases, the password in an MQCSP structure for a client application will be sent across a network in plain text. To ensure that client application passwords are protected appropriately, see IBM MQCSP password protection.

Here is an example of an xa\_open string:

```
channel=MARS.SVR,trptype=tcp,connname=MARS(1415),qmname=MARS,tpm=cics
```


## The CHANNEL, TRPTYPE, CONNAME, and QMNAME parameters of the xa\_open string:


Use this information to understand how the extended transactional client uses these parameters to determine the queue manager to connect to.


If the CHANNEL and CONNAME parameters are supplied in the xa\_open string, the extended transactional client uses these parameters and the TRPTYPE parameter to start an MQI channel to the server queue manager.

If the CHANNEL and CONNAME parameters are not supplied in the xa\_open string, the extended transactional client uses the value of the MQSERVER environment variable to start an MQI channel. If the MQSERVER environment variable is not defined, the extended transactional client uses the entry in the client channel definition identified by the QMNAME parameter.

In each of these cases, the extended transactional client checks that the value of the QMNAME parameter is the name of the queue manager at the server end of the MQI channel. If it is not, the xa\_open call fails and the transaction manager reports the failure to the application.

If the application connects to a queue manager at an earlier version than V7.0.1, the xa\_open call succeeds but the transaction has a QMGR unit of recovery disposition.  Ensure that applications that require the GROUP unit of recovery disposition connect only to queue managers at V7.0.1 or later.

 If the application uses a QSG name in QMNAME parameter field and the GROUPUR property is disabled on the queue manager to which it connects then the xa\_open call fails.

 If the applicat.ibmion client is connecting to a z/OS queue manager at V7.0.1 or later it can specify a queue-sharing group (QSG) name for the QMNAME parameter. This allows the application client to participate in a transaction with a GROUP unit of recovery disposition. For more information about the GROUP unit of recovery disposition, see Unit of recovery disposition.

When the client application later calls MQCONN or MQCONNX on the same thread that the transaction manager used to issue the xa\_open call, the application receives a connection handle for the MQI channel that was started by the xa\_open call. A second MQI channel is not started. The extended transactional client checks that the value of the *QMgrName* parameter on the MQCONN or MQCONNX call is the name of the queue manager at the server end of the MQI channel. If it is not, the MQCONN or MQCONNX call fails with a reason code of MQRC\_ANOTHER\_Q\_MGR\_CONNECTED. If the value of the *QMgrName* parameter is blank or a single asterisk (\*), or starts with an asterisk, the MQCONN or MQCONNX call fails with a reason code of MQRC\_Q\_MGR\_NAME\_ERROR.

If the client application has already started an MQI channel by calling MQCONN or MQCONNX before the transaction manager calls xa\_open on the same thread, the transaction manager uses this MQI channel instead. A second MQI channel is not started. The extended transactional client checks that the value of the QMNAME parameter in the xa\_open string is the name of the server queue manager. If it is not, the xa\_open call fails.

If a client application starts an MQI channel first, the value of the *QMgrName* parameter on the MQCONN or MQCONNX call can be blank or a single asterisk (\*), or it can start with an asterisk. Under these circumstances, however, you must ensure that the queue manager to which the application connects is the same as the queue manager that the transaction manager intends to open as a resource manager when it later calls xa\_open on the same thread. You might encounter fewer problems, therefore, if the value of the *QMgrName* parameter identifies the queue manager explicitly by name.

### The TPM and AXLIB parameters:

An extended transactional client uses the TPM and AXLIB parameters to locate the transaction manager's ax\_reg and ax\_unreg functions. These functions are used only if the queue manager uses dynamic registration.

If the TPM parameter is supplied in an xa\_open string, but the AXLIB parameter is not supplied, the extended transactional client assumes a value for the AXLIB parameter based on the value of the TPM parameter. See Table 113 for the assumed values of the AXLIB parameter.

Table 113. Assumed values of the AXLIB parameter

| Value of TPM | Platform        | Assumed value of AXLIB                              |
|--------------|-----------------|-----------------------------------------------------|
| CICS         | AIX             | /usr/lpp/encina/lib/libEncServer.a(EncServer_shr.o) |
| CICS         | HP-UX           | /opt/encina/lib/libEncServer.sl                     |
| CICS         | Solaris         | /opt/encina/lib/libEncServer.so                     |
| CICS         | Windows systems | libEncServer                                        |
| Tuxedo       | AIX             | /usr/lpp/tuxedo/lib/libtux.a(libtux.so.60)          |
| Tuxedo       | HP-UX           | /opt/tuxedo/lib/libtux.sl                           |
| Tuxedo       | Solaris         | /opt/tuxedo/lib/libtux.so.60                        |
| Tuxedo       | Windows systems | libtux                                              |

If the AXLIB parameter is supplied in an xa\_open string, the extended transactional client uses its value to override any assumed value based on the value of the TPM parameter. The AXLIB parameter can also be used for a transaction manager for which the TPM parameter does not have a specified value.

### Additional error processing for xa\_open:

The xa\_open call fails in certain circumstances.

Topics in this section describe situations in which the xa\_open call fails. It also fails if any of the following situations occur:

- There are errors in the xa\_open string.
- There is insufficient information to start an MQI channel.
- There is a problem while trying to start an MQI channel (the server queue manager is not running, for example).

## Recovery following a failure in extended transactional processing:

Following a failure, a transaction manager must be able to recover any incomplete units of work. To do this, the transaction manager must be able to open as a resource manager any queue manager that was participating in an incomplete unit of work at the time of the failure.

Therefore, you must ensure that all incomplete units of work have been resolved before making changes to any configuration information.

Alternatively, you must ensure that the configuration changes do not affect the ability of the transaction manager to open the queue managers it needs to open. The following are examples of such configuration changes:

- Changing the contents of an xa\_open string
- Changing the value of the MQSERVER environment variable
- Changing entries in the client channel definition table (CCDT)
- Deleting a server connection channel definition

## The XA switch structures:

Two XA switch structures are supplied with the extended transactional client on each platform.

These switch structures are:

### **MQRMIXASwitch**

This switch structure is used by a transaction manager when a queue manager, acting as a resource manager, is not using dynamic registration.

### **MQRMIXASwitchDynamic**

This switch structure is used by a transaction manager when a queue manager, acting as a resource manager, uses dynamic registration.

These switch structures are located in the libraries shown in Table 114.

Table 114. IBM MQ libraries containing the XA switch structures

| Platform                         | Library containing the XA switch structures             |
|----------------------------------|---------------------------------------------------------|
| AIX<br>HP-UX<br>Linux<br>Solaris | <i>MQ_INSTALLATION_PATH</i> /lib/libmqcxa               |
| Windows systems                  | <i>MQ_INSTALLATION_PATH</i> \bin\mqcxa.dll <sup>1</sup> |

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

The name of the IBM MQ resource manager in each switch structure is MQSeries\_XA\_RMI, but many queue managers can share the same switch structure.

**Related concepts:**

“Dynamic registration and extended transactional processing”

Using dynamic registration is a form of optimization because it can reduce the number of xa\_ function calls issued by the transaction manager.

*Dynamic registration and extended transactional processing:*

Using dynamic registration is a form of optimization because it can reduce the number of xa\_ function calls issued by the transaction manager.

If a queue manager does not use dynamic registration, a transaction manager involves the queue manager in every unit of work. The transaction manager does this by calling xa\_start, xa\_end, and xa\_prepare, even if the queue manager has no resources that are updated within the unit of work.

If a queue manager uses dynamic registration, a transaction manager starts by assuming that the queue manager is not involved in a unit of work, and does not call xa\_start. The queue manager then becomes involved in the unit of work only if its resources are updated within sync point control. If this occurs, the extended transactional client calls ax\_reg to register the queue manager's involvement.

**Using the extended transactional client with SSL channels:**

You cannot set up an SSL channel using the xa\_open string. Follow these instructions to use the client channel definition table (ccdt).

**About this task**

Because of the limited size of the xa\_open xa\_info string, it is not possible to pass all the information required to set up an SSL channel using the xa\_open string method of connecting to a queue manager. Therefore you must either use the client channel definition table or, if your transaction manager allows, create the channel with MQCONN before issuing the xa\_open call.

To use the client channel definition table, follow these steps:

**Procedure**

1. Specify an xa\_open string containing only the mandatory qmname (queue manager name) parameter, for example: XA\_Open\_String=qmname=MYQM
2. Use a queue manager to define a CLNTCONN (client-connection) channel with the required SSL parameters. Include the queue manager name in the QMNAME attribute on the CLNTCONN definition. This will be matched up with the qmname in the xa\_open string.
3. Make the CLNTCONN definition available to the client system in a client channel definition table (CCDT) or, on Windows, in the active directory.
4. If you are using a CCDT, identify the CCDT containing the definition of the CLNTCONN channel using environment variables MQCHLLIB and MQCHLTAB. Set these variables in the environments used by both the client application and the transaction manager.

**Results**

This gives the transaction manager a channel definition to the appropriate queue manager with the SSL attributes needed to authenticate correctly, including SSLCIPH, the CipherSpec.

## Configuring an extended transactional client for CICS:

You configure an extended transactional client for use by CICS by adding an XAD resource definition to a CICS region.

Add the XAD resource definition by using the CICS resource definition online (RDO) command, **cicsadd**. The XAD resource definition specifies the following information:

- An xa\_open string
- The fully qualified path name of a switch load file

One switch load file is supplied for use by CICS on each of the following platforms: AIX, HP-UX, Solaris, and Windows systems. Each switch load file contains a function that returns a pointer to the XA switch structure that is used for dynamic registration, MQRMIXASwitchDynamic. See Table 115 for the fully qualified path name of each switch load file.

Table 115. The switch load files

| Platform                         | Switch load file                                                |
|----------------------------------|-----------------------------------------------------------------|
| AIX<br>HP-UX<br>Linux<br>Solaris | <code>MQ_INSTALLATION_PATH/lib/amqczsc</code>                   |
| Windows systems                  | <code>MQ_INSTALLATION_PATH\bin\mqcc4swi.dll</code> <sup>1</sup> |

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Here is an example of an XAD resource definition for Windows systems:

```
cicsadd -c xad -r REGION1 WMQXA \
 ResourceDescription="IBM MQ queue manager MARS" \
 XAOpen="channel=MARS.SVR,trptype=tcp,connname=MARS(1415),qmname=MARS,tpm=cics" \
 SwitchLoadFile="C:\Program Files\IBM\WebSphere MQ\bin\mqcc4swi.dll"
```

For more information about adding an XAD resource definition to a CICS region, see the *CICS Administration Reference* and the *CICS Administration Guide* for your platform.

Note the following information about using CICS with an extended transactional client:

- You can add only one XAD resource definition for IBM MQ to a CICS region. This means that only one queue manager can be associated with a region, and all CICS applications that run in the region can connect only to that queue manager. If you want to run CICS applications that connect to a different queue manager, you must run the applications in a different region.
- Each application server in a region calls xa\_open while it is initializing and starts an MQI channel to the queue manager associated with the region. This means that the queue manager must be started before an application server starts, otherwise the xa\_open call fails. All IBM MQ MQI client applications later processed by the application server use the same MQI channel.
- When an MQI channel starts, and there is no security exit at the client end of the channel, the user ID that flows from the client system to the server connection MCA is cics. Under certain circumstances, the queue manager uses this user ID for authority checks when the server connection MCA subsequently attempts to access the queue manager resources on behalf of a client application. If this user ID is used for authority checks, you must ensure that it has the authority to access all the resources it needs to access.

For information about when the queue manager uses this user ID for authority checks, see Security.

- The CICS task termination exits that are supplied for use on IBM MQ client systems are listed in Table 116 on page 711. You configure these exits in the same way that you configure the corresponding exits for IBM MQ server systems. For this information, therefore, see the Enabling CICS user exits.

Table 116. CICS task termination exits

| Platform                         | Source     | Library      |
|----------------------------------|------------|--------------|
| AIX<br>HP-UX<br>Linux<br>Solaris | amqzscgx.c | amqczscg     |
| Windows systems                  | amqzscgn.c | mqqc1415.dll |

### Configuring an extended transactional client for Tuxedo:

To configure XAD resource definition for use by Tuxedo, update the UBBCONFIG file and resource manager table.

To configure XAD resource definition for use by Tuxedo, perform the following actions:

- In the GROUPS section of the Tuxedo UBBCONFIG file for an application, use the OPENINFO parameter to specify an xa\_open string.

For an example of how to do this, see the sample UBBCONFIG file, which is supplied for use with the Tuxedo sample programs. On AIX, HP-UX, and Solaris, the name of the file is ubbstxcx.cfg and, on Windows systems, the name of the file is ubbstxcn.cfg.

- In the entry for a queue manager in the Tuxedo resource manager table:
  - udataobj/RM ( AIX, HP-UX, and Solaris)
  - udataobj\rm ( Windows systems)

specify the name of an XA switch structure and the fully qualified path name of the library that contains the structure. For an example of how to do this for each platform, see TUXEDO samples. Tuxedo supports dynamic registration of a resource manager, and so you can use either MQRMIXASwitch or MQRMIXASwitchDynamic.

### Microsoft Transaction Server

No additional configuration is required before you can use MTS as a transaction manager. However, there are some points to note.

Note the following information about using MTS with the extended transactional client:

- An MTS application always starts an MQI channel when it connects to a server queue manager. MTS, in its role as a transaction manager, then uses the same MQI channel to communicate with the queue manager.
- Following a failure, MTS must be able to recover any incomplete units of work. To do this, MTS must be able to communicate with any queue manager that was participating in an incomplete unit of work at the time of the failure.

When an MTS application connects to a server queue manager and starts an MQI channel, the extended transactional client extracts sufficient information from the MQCONN or MQCONNX call to enable the channel to be restarted following a failure, if required. The extended transactional client passes the information to MTS, and MTS records the information in its log.

If the MTS application issues an MQCONN call, this information is simply the name of the queue manager. If the MTS application issues an MQCONNX call and provides a channel definition structure, MQCD, the information also includes the name of the MQI channel, the network address of the server queue manager, and the communications protocol for the channel.

In a recovery situation, MTS passes this information back to the extended transactional client, and the extended transactional client uses it to restart the MQI channel.

If you ever need to change any configuration information, therefore, ensure that all incomplete units of work have been resolved before making the changes. Alternatively, ensure that the configuration

changes do not affect the ability of the extended transactional client to restart an MQI channel using the information recorded by MTS. The following are examples of such configuration changes:

- Changing the value of the MQSERVER environment variable
- Changing entries in the client channel definition table (CCDT)
- Deleting a server connection channel definition
- Note the following conditions when using an extended transactional client with MTS:
  - Within a single thread, a client application can be connected to only one queue manager at a time.
  - Each thread of a client application can connect to a different queue manager.
  - A client application cannot use shared connection handles.

## Defining MQI channels

To create a new channel, you have to create **two** channel definitions, one for each end of the connection, using the same channel name and compatible channel types. In this case, the channel types are *server-connection* and *client-connection*.

### User defined channels

When the server does not automatically define channels there are two ways of creating the channel definitions and giving the IBM MQ application on the IBM MQ MQI client machine access to the channel.

These two methods are described in detail:

1. Create one channel definition on the IBM MQ client and the other on the server.

This applies to any combination of IBM MQ MQI client and server platforms. Use it when you are getting started on the system, or to test your setup.

See “Creating server-connection and client-connection definitions on different platforms” on page 713 for details on how to use this method.

2. Create both channel definitions on the server machine.

Use this method when you are setting up multiple channels and IBM MQ MQI client machines at the same time.

See “Creating server-connection and client-connection definitions on the server” on page 716 for details on how to use this method.

### Automatically defined channels

IBM MQ products on platforms other than z/OS include a feature that can automatically create a channel definition on the server if one does not exist.

If an inbound attach request is received from a client and an appropriate server-connection definition cannot be found on that queue manager, IBM MQ creates a definition automatically and adds it to the queue manager. The automatic definition is based on the definition of the default server-connection channel SYSTEM.AUTO.SVRCONN. You enable automatic definition of server-connection definitions by updating the queue manager object using the ALTER QMGR command with the CHAD parameter (or the PCF command Change Queue Manager with the ChannelAutoDef parameter).



### Related concepts:



“Channel control function” on page 832

The channel control function provides facilities for you to define, monitor, and control channels.

## Creating server-connection and client-connection definitions on different platforms

You can create each channel definition on the computer to which it applies. There are restrictions on how you can create channel definitions on a client computer.

On all platforms, you can use IBM MQ Script (MQSC) commands, programmable command format (PCF) commands, or the IBM MQ Explorer to define a server-connection channel on the server machine.

 On z/OS you can also use the Operation and Control panels.  On IBM i you can also use the panel interface.

Because MQSC commands are not available on a machine where IBM MQ has been installed as an IBM MQ MQI client only, you must use different ways of defining a client-connection channel on the client machine.

### Using runmqsc

You can specify the **-c** parameter and, optionally, the **-u** parameter to connect **runmqsc** as a client to the queue manager you want to administer.

If you use the **-u** parameter to supply a user ID, you are prompted for a matching password.

If you have configured the CONNAUTH AUTHINFO record with CHCKLOCL(REQUIRED) or CHCKLOCL(REQDADM), you must use the **-u** parameter otherwise you will not be able to administer your queue manager with **runmqsc**.

### Related concepts:

“Creating a client-connection channel on the IBM MQ MQI client” on page 714

You can define a client-connection channel on the client workstation using MQSERVER or using the MQCNO structure on an MQCONN call.

### Related tasks:

“Defining a server-connection channel on the server”

Start MQSC if necessary, then define the server-connection channel.

## Defining a server-connection channel on the server

Start MQSC if necessary, then define the server-connection channel.

### Procedure

- Optional: If your server platform is not z/OS, first create and start a queue manager and then start MQSC commands.
  - Create a queue manager, called QM1 for example:

```
crtmqm QM1
```
  - Start the queue manager:

```
strmqm QM1
```
  - Start MQSC commands:

```
runmqsc QM1
```
- Define a channel with your chosen name and a channel type of *server-connection*.

```
DEFINE CHANNEL(CHAN1) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Server-connection to Client_1')
```

This channel definition is associated with the queue manager running on the server.

3. Use the following command to allow the inbound connect access to your queue manager:

```
SET CHLAUTH(CHAN1) TYPE(ADDRESSMAP) ADDRESS('IP address') MCAUSER('userid')
```

- Where SET CHLAUTH uses the name of the channel defined in the previous step.
- Where 'IP address' is the IP address of the client.
- Where 'userid' is the ID you want to provide to the channel for access control to the target queues. This field is case-sensitive.

You can choose to identify your inbound connection using a number of different attributes. The example uses IP address. Alternative attributes include client user ID and SSL or TLS Subject Distinguished Name. For more information, see Channel authentication records

## Creating a client-connection channel on the IBM MQ MQI client

You can define a client-connection channel on the client workstation using MQSERVER or using the MQCNO structure on an MQCONN call.

### Using MQSERVER

You can use the MQSERVER environment variable to specify a simple definition of a client-connection channel. It is simple in the sense that you can specify only a few attributes of the channel using this method.

- Specify a simple channel definition on Windows as follows:  
`SET MQSERVER=ChannelName/TransportType/ConnectionName`
- Specify a simple channel definition on UNIX and Linux systems as follows:  
`export MQSERVER=ChannelName/TransportType/ConnectionName`
- Specify a simple channel definition on IBM i systems as follows:  
`ADDENVVAR ENVVAR(MQSERVER) VALUE('ChannelName/TransportType/ConnectionName')`

where:

- ChannelName must be the same name as defined on the server. It cannot contain a forward slash.
- TransportType can be one of the following values, depending on your IBM MQ MQI client platform:
  - LU62
  - TCP
  - NETBIOS
  - SPX

**Note:** On UNIX and Linux systems, the TransportType is case-sensitive and must be uppercase. An MQCONN or MQCONN call returns 2058 if the TransportType is not recognized

- ConnectionName is the name of the server as defined to the communications protocol (TransportType).

For example, on Windows:

```
SET MQSERVER=CHANNEL1/TCP/MCID66499
```

or, on UNIX and Linux systems:

```
export MQSERVER=CHANNEL1/TCP/'MCID66499'
```

**Note:** To change the TCP/IP port number, see "MQSERVER" on page 751.





Figure 78. Simple channel definition

Some more examples of simple channel definitions are:

- On Windows:
 

```
SET MQSERVER=CHANNEL1/TCP/9.20.4.56
SET MQSERVER=CHANNEL1/NETBIOS/BOX643
```
- On UNIX and Linux systems:
 

```
export MQSERVER=CHANNEL1/TCP/'9.20.4.56'
export MQSERVER=CHANNEL1/LU62/BOX99
```
-  On IBM i:
 

```
ADDENVVAR ENVVAR(MQSERVER) VALUE('CHANNEL1/TCP/9.20.4.56(1416)')
```

where BOX99 is the LU 6.2 ConnectionName.

On the IBM MQ MQI client, all **MQCONN** or **MQCONNX** requests then attempt to use the channel you have defined, unless the channel is overridden in an **MQCD** structure referenced from the **MQCNO** structure supplied to **MQCONNX**.

**Note:** For more information about the *MQSERVER* environment variable, see “MQSERVER” on page 751.

### Using the MQCNO structure on an MQCONNX call

An IBM MQ MQI client application can use the connect options structure, **MQCNO**, on an **MQCONNX** call to reference a channel definition structure, **MQCD**, that contains the definition of a client-connection channel.

In this way, the client application can specify the **ChannelName**, **TransportType**, and **ConnectionName** attributes of a channel at run time, enabling the client application to connect to multiple server queue managers simultaneously.

Note that if you define a channel using the *MQSERVER* environment variable, it is not possible to specify the **ChannelName**, **TransportType**, and **ConnectionName** attributes at run time.

A client application can also specify attributes of a channel such as **MaxMsgLength** and **SecurityExit**. Specifying such attributes enables the client application to specify values for the attributes that are not the default values, and enables channel exit programs to be called at the client end of an MQI channel.

If a channel uses the Secure Sockets Layer (SSL) or Transport Layer Security (TLS), a client application can also provide information relating to SSL or TLS in the **MQCD** structure. Additional information



relating to SSL or TLS can be provided in the SSL or TLS configuration options structure, MQSCO, which is also referenced by the MQCNO structure on an **MQCONN** call.

For more information about the MQCNO, MQCD, and MQSCO structures, see MQCNO, MQCD, and MQSCO.

**Note:** The sample program for MQCONN is called **amqscnxc**. Another sample program called **amqssl1c** demonstrates use of the MQSCO structure.

## Creating server-connection and client-connection definitions on the server

You can create both definitions on the server, then make the client-connection definition available to the client.

First define a server-connection channel and then define a client-connection channel. On all platforms, you can use IBM MQ Script (MQSC) commands, programmable command format (PCF) commands or the IBM MQ Explorer to define a server-connection channel on the server machine.  On z/OS you can also use the Operation and Control panels.  On IBM i you can also use the panel interface.

Client-connection channel definitions created on the server are made available to clients using a client channel definition table (CCDT).

### Related concepts:

“Client channel definition table”

The client channel definition table (CCDT) determines the channel definitions and authentication information used by client applications to connect to the queue manager. On platforms other than z/OS a CCDT is created automatically. You must then make it available to the client application.

### Related tasks:

“Defining the server-connection channel on the server” on page 719

Create a server-connection channel definition for the queue manager.

“Defining the client-connection channel on the server” on page 719

Having defined the server-connection channel, you now define the corresponding client-connection channel.

“Accessing client-connection channel definitions” on page 720

Make the client channel definition table (CCDT) available to client applications by copying or sharing it, then specify its location and name on the client computer.

## Client channel definition table

The client channel definition table (CCDT) determines the channel definitions and authentication information used by client applications to connect to the queue manager. On platforms other than z/OS a CCDT is created automatically. You must then make it available to the client application.

The purpose of the client channel definition table (CCDT) is to determine the channel definitions used by client applications to connect to the queue manager. The channel definition also specifies the authentication information that applies to the connections.

The CCDT is a binary file. It is generated by a queue manager. The queue manager does not read the CCDT file.

On platforms other than z/OS, the CCDT is created when the queue manager is created. Client connection channels are added to the table when you use the **DEFINE CHANNEL** command, and their definitions altered when you issue the **ALTER CHANNEL** command.

You can use the CCDT to provide clients with the authentication information to check for SSL certificate revocation. Define a namelist containing authentication information objects and set the queue manager attribute **SSLCRLNameList** to the name of the namelist.




There are a number of ways for a client application to use a CCDT. The CCDT can be copied to the client computer. You can copy the CCDT to a location shared by more than one client. You can make the CCDT accessible to the client as a shared file, while it remains located on the server.

If you use FTP to copy the file, use the **bin** option to set binary mode; do not use the default ASCII mode. Whichever method you choose to make the CCDT available, the location must be secure to prevent unauthorized changes to the channels.


## Server platforms other than z/OS

A default CCDT called **AMQCLCHL.TAB** is created when you create a queue manager.

By default, **AMQCLCHL.TAB** is located in the following directory on a server:

-  On IBM i, in the integrated file system:  
`/QIBM/UserData/mqm/qmgrs/QUEUEMANAGERNAME/&ipcc`
-   On UNIX and Linux systems:  
`/prefix/qmgrs/QUEUEMANAGERNAME/@ipcc`

The name of the directory referenced by *QUEUEMANAGERNAME* is case-sensitive on UNIX and Linux systems. The directory name might not be the same as the queue manager name, if the queue manager name has special characters in it.

-  On Windows:  
`MQ_INSTALLATION_PATH\data\qmgrs\QUEUEMANAGERNAME\@ipcc`

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

However, you might have chosen to use a different directory for queue manager data. You can specify the parameter **-md** *DataPath* when you used the **crtmqm** command. If you do, **AMQCLCHL.TAB** is located in the *@ipcc* directory of the *DataPath* you specified.

The path to the CCDT can be changed by setting **MQCHLLIB**. If you do set **MQCHLLIB**, be aware, if you have multiple queue managers on the same server, they share the same CCDT location.

The CCDT is created when the queue manager is created. Each entry of a CCDT represents a client connection to a specific queue manager. A new entry is added when you define a client-connection channel using the **DEFINE CHANNEL** command, and the entry is updated when you alter the client-connection channels by using the **ALTER CHANNEL** command.

## Client platforms at IBM MQ Version 8.0

You can create a CCDT on the client machine directly by using the **runmqsc** command with the **-n** parameter. The CCDT will be created in the location indicated by **MQCHLLIB** and with the filename indicated by **MQCHLTAB** which is **AMQCLCHL.TAB** by default.

Note, that if you specify the **-n** parameter, you must not specify any other parameter.

Each entry of a CCDT represents a client connection to a specific queue manager. A new entry is added when you define a client-connection channel using the **DEFINE CHANNEL** command, and the entry is updated when you alter the client-connection channels by using the **ALTER CHANNEL** command.

## How to specify the location of the CCDT on the client

On a client system, you can specify the location of the CCDT in two ways:

- Using the environment variables MQCHLLIB to specify the directory where the table is located, and MQCHLTAB to specify the file name of the table.
- Using the client configuration file. In the CHANNELS stanza, use the attributes ChannelDefinitionDirectory to specify the directory where the table is located, and ChannelDefinitionFile to specify the file name.

If the location is specified both in the client configuration file and by using environment variables, the environment variables take priority. You can use this feature to specify a standard location in the client configuration file and override it using environment variables when necessary.

### Related reference:

“MQCHLLIB” on page 749

MQCHLLIB specifies the directory path to the file containing the client channel definition table (CCDT). The file is created on the server, but can be copied across to the IBM MQ MQI client workstation.

### Related information:

Working with revoked certificates

### Migration and client channel definition tables (CCDT):

In general, the internal format of the client channel definition table might change from one release level of IBM MQ to the next. As a result, an IBM MQ MQI client can use a client channel definition table only when it has been prepared by a server queue manager that is at the same release level as the client, or at an earlier release level.

A Version 7.1 IBM MQ MQI client can use a client channel definition table that has been prepared by a Version 6.0 queue manager. But a Version 6.0 client cannot use a client channel definition table that has been prepared by a Version 7.1 queue manager.

## Client connection channels in the Active Directory

On Windows systems that support the Active Directory, IBM MQ publishes client connection channels in the Active Directory to provide dynamic client-server binding.

When client connection channel objects are defined, they are written into a client channel definition file, called AMQCLCHL.TAB by default. If the client connection channels use the TCP/IP protocol, the IBM MQ server also publishes them in the Active Directory. When the IBM MQ client determines how to connect to the server, it looks for a relevant client connection channel object definition using the following search order:

1. MQCONNX MQCD data structure
2. MQSERVER environment variable
3. client channel definition file
4. Active Directory

This order means that any current applications are not affected by any change. You can think of these entries in the Active Directory as records in the client channel definition file, and the IBM MQ client processes them in the same way. To configure and administer support for publishing client connection channel definitions in the Active Directory, use the **setmqscp** command, as described in **setmqscp**.

## Defining the server-connection channel on the server

Create a server-connection channel definition for the queue manager.

### Procedure

1. On the server machine, define a channel with your chosen name and a channel type of *server-connection*. For example:  

```
DEFINE CHANNEL(CHAN2) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Server-connection to Client_2')
```
2. Use the following command to allow the inbound connect access to your queue manager:  

```
SET CHLAUTH(CHAN2) TYPE(ADDRESSMAP) ADDRESS('IP address') MCAUSER('userid')
```

  - Where SET CHLAUTH uses the name of the channel defined in the previous step.
  - Where 'IP address' IP address is the IP address of the client.
  - Where 'userid' is the ID you want to provide to the channel for access control to the target queues. This field is case-sensitive.

You can choose to identify your inbound connection using a number of different attributes. The example uses IP address. Alternative attributes include client user ID and SSL or TLS Subject Distinguished Name. For more information, see [Channel authentication records](#). This channel definition is associated with the queue manager running on the server.




Figure 79. Defining the server-connection channel

## Defining the client-connection channel on the server

Having defined the server-connection channel, you now define the corresponding client-connection channel.

### Before you begin

Define the server-connection channel.

### Procedure

1. Define a channel with the same name as the server-connection channel, but a channel type of *client-connection*. You must state the connection name (CONNNAME). For TCP/IP, the connection name is the network address or host name of the server machine. It is also advisable to specify the queue manager name (QMNAME) to which you want your IBM MQ application, running in the client environment, to connect. By varying the queue manager name, you can define a set of channels to connect to different queue managers.  

```
DEFINE CHANNEL(CHAN2) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(9.20.4.26) QMNAME(QM2) DESCR('Client-connection to Server_2')
```
2. Use the following command to allow the inbound connect access to your queue manager:  

```
SET CHLAUTH(CHAN2) TYPE(ADDRESSMAP) ADDRESS('IP-address') MCAUSER('userid')
```

- Where SET CHLAUTH uses the name of the channel defined in the previous step.
- Where 'IP address' is the IP address of the client.
- Where 'userid' is the ID you want to provide to the channel for access control to the target queues. This field is case-sensitive.

You can choose to identify your inbound connection using a number of different attributes. The example uses IP address. Alternative attributes include client user ID and SSL or TLS Subject Distinguished Name. For more information, see Channel authentication records

## Results

On platforms other than z/OS, this channel definition is stored in a file called the client channel definition table (CCDT), which is associated with the queue manager. The client channel definition table can contain more than one client-connection channel definition. For more information about the client channel definition table, and for the corresponding information about how client-connection channel definitions are stored on z/OS, see "Client channel definition table" on page 716.




Figure 80. Defining the client-connection channel

## Accessing client-connection channel definitions

Make the client channel definition table (CCDT) available to client applications by copying or sharing it, then specify its location and name on the client computer.

### Before you begin

You have defined the client-connection channels you need.

On z/OS, you have created a CCDT. On other platforms, the CCDT is automatically created and updated.

### About this task

For a client application to use the client channel definition table (CCDT), you must make the CCDT available to it and specify its location and name

### Procedure

1. Make the CCDT available to the client applications in one of three ways:
  - a. Optional: Copy the CCDT to the client computer.
  - b. Optional: Copy the CCDT to a location shared by more than one client.



- c. Optional: Leave the CCDT on the server but make it shareable by the client.

Whichever location you choose for the CCDT, the location must be secure to prevent unauthorized changes to the channels.

2. On the client, specify the location and name of the file containing the CCDT in one of three ways:
  - a. Optional: Use the CHANNELS stanza of the client configuration file. For more information, see “CHANNELS stanza of the client configuration file” on page 734.
  - b. Optional: Use the environment variables MQCHLLIB and MQCHLTAB.

For example, you can set the environment variables by typing:

- On HP Integrity NonStop Server, and UNIX and Linux systems:

```
export MQCHLLIB=MQ_INSTALLATION_PATH/qmgrs/QUEUEMANAGERNAME/@ipcc
export MQCHLTAB=AMQCLCHL.TAB
```

-  On IBM i:

```
ADDENVVAR ENVVAR(MQCHLLIB) VALUE('/QIBM/UserData/mqm/qmgrs/QUEUEMANAGERNAME/@ipcc')
ADDENVVAR ENVVAR(MQCHLTAB) VALUE(AMQCLCHL.TAB)
```

where *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

- c. Optional: On Windows only, use the **setmqscp** control command to publish the client-connection channel definitions in Active Directory

If the MQSERVER environment variable is set, an IBM MQ client uses the client-connection channel definition specified by MQSERVER in preference to any definitions in the client channel definition table.

## Channel-exit programs for MQI channels

Three types of channel exit are available to the IBM MQ MQI client environment on UNIX, Linux and Windows systems.

These are:

- Send exit
- Receive exit
- Security exit

These exits are available at both the client and the server end of the channel. Exits are not available to your application if you are using the MQSERVER environment variable. Channel exits are explained in Channel exit programs for messaging channels.

The send and receive exits work together. There are several possible ways in which you can use them:

- Splitting and reassembling a message
- Compressing and decompressing data in a message (this functionality is provided as part of IBM MQ, but you might want to use a different compression technique)
- Encrypting and decrypting user data (this functionality is provided as part of IBM MQ, but you might want to use a different encryption technique)
- Journaling each message sent and received

You can use the security exit to ensure that the IBM MQ client and server are correctly identified, and to control access.

If send or receive exits on the server-connection side of the channel instance need to perform MQI calls on the connection with which they are associated, they use the connection handle provided in the MQCXP Hconn field. You must be aware that client-connection send and receive exits cannot make MQI calls.

**Related concepts:**

“Security exits on a client connection”

You can use security exit programs to verify that the partner at the other end of a channel is genuine. Special considerations apply when a security exit is applied to a client connection.

**Related reference:**

“Path to exits”

A default path for location of the channel exits is defined in the client configuration file. Channel exits are loaded when a channel is initialized.

“Identifying the API call in a send or receive exit program” on page 724

When you use MQI channels for clients, byte 10 of the agent buffer identifies the API call in use when a send or receive exit is called. This is useful for identifying which channel flows include user data and might require processing such as encryption or digital signing.

**Related information:**

Extending queue manager facilities

User exits, API exits, and IBM MQ installable services

**Path to exits**

A default path for location of the channel exits is defined in the client configuration file. Channel exits are loaded when a channel is initialized.

On UNIX, Linux and Windows systems, a client configuration file is added to your system during installation of the IBM MQ MQI client. A default path for location of the channel exits on the client is defined in this file, using the stanza:

```
ClientExitPath:
ExitsDefaultPath= string
ExitsDefaultPath64= string
```

where *string* is a file location in a format appropriate to the platform

When a channel is initialized, after an **MQCONN** or **MQCONNX** call, the client configuration file is searched. The ClientExitPath stanza is read and any channel exits that are specified in the channel definition are loaded.

**Security exits on a client connection**

You can use security exit programs to verify that the partner at the other end of a channel is genuine. Special considerations apply when a security exit is applied to a client connection.

Figure 81 on page 724 illustrates the use of security exits in a client connection, using the IBM MQ object authority manager to authenticate a user. Either SecurityParmsPtr or SecurityParmsOffset is set in the MQCNO structure on the client and there are security exits at both ends of the channel. After the normal security message exchange has ended, and the channel is ready to run, the MQCSP structure accessed from the MQCXP SecurityParms field is passed to the security exit on the client. The exit type is set to MQXR\_SEC\_PARMS. The security exit can elect to do nothing to the user identifier and password, or it can alter either or both of them. The data returned from the exit is then sent to the server-connection end of the channel. The MQCSP structure is rebuilt on the server-connection end of the channel and is passed to the server-connection security exit accessed from the MQCXP SecurityParms field. The security exit receives and processes this data. This processing is typically to reverse any change made to the user ID and password fields in the client exit, which are then used to authorize the queue manager connection. The resulting MQCSP structure is referenced using SecurityParmsPtr in the MQCNO structure on the queue manager system.

The memory address that is passed back by the MQCXP SecurityParms field must remain addressable and unchanged until MQXR\_TERM. An exit must not invalidate or free the memory back to the system before the exit is called for MQXR\_TERM.

If `SecurityParmsPtr` or `SecurityParmsOffset` are set in the `MQCNO` structure and there is a security exit at only one end of the channel, the security exit receives and processes the `MQCSP` structure. Actions such as encryption are inappropriate for a single user exit, as there is no exit to perform the complementary action.

If `SecurityParmsPtr` and `SecurityParmsOffset` are not set in the `MQCNO` structure and there is a security exit at either or both ends of the channel, the security exit or exits are called. Either security exit can return its own `MQCSP` structure, addressed through the `SecurityParmsPtr`; the security exit is not called again until it is terminated (`ExitReason` of `MQXR_TERM`). The exit writer can free the memory used for the `MQCSP` at that stage.

When a server-connection channel instance is sharing more than one conversation, the pattern of calls to the security exit is restricted on the second and subsequent conversations.

For the first conversation, the pattern is the same as if the channel instance is not sharing conversations. For the second and subsequent conversations, the security exit is never called with `MQXR_INIT`, `MQXR_INIT_SEC`, or `MQXR_SEC_MSG`. It is called with `MQXR_SEC_PARMS`.

In a channel instance with sharing conversations, `MQXR_TERM` is called only for the last conversation running.

Each conversation has the opportunity in the `MQXR_SEC_PARMS` invocation of the exit to alter the `MQCD`; on the server-connection end of the channel this feature can be useful to vary, for example, the `MCAUserIdentifier` or `LongMCAUserIdPtr` values before the connection is made to the queue manager.

| Server-connection exit                                          | Client-connection exit                                          |
|-----------------------------------------------------------------|-----------------------------------------------------------------|
|                                                                 | Invoked with MQXR_INIT<br>Responds with MQXCC_OK                |
| Invoked with MQXR_INIT<br>Responds with MQXCC_OK                |                                                                 |
|                                                                 | Invoked with MQXR_INIT_SEC<br>Responds with MQXCC_OK            |
| Invoked with MQXR_INIT_SEC<br>Responds with MQXCC_OK            |                                                                 |
|                                                                 | Invoked with MQXR_SEC_PARMS<br>Responds with MQXCC_SEC_SEND_MSG |
| Invoked with MQXR_SEC_PARMS<br>Responds with MQXCC_SEC_SEND_MSG |                                                                 |
| Data transfer begins                                            |                                                                 |
| Invoked with MQXR_TERM<br>Responds with MQXCC_OK                | Invoked with MQXR_TERM<br>Responds with MQXCC_OK                |

Figure 81. Client connection-initiated exchange with agreement for client connection using security parameters

**Note:** Security exit applications constructed prior to the release of IBM MQ v7.1 may require updating. For more information see Channel security exit programs.

### Identifying the API call in a send or receive exit program

When you use MQI channels for clients, byte 10 of the agent buffer identifies the API call in use when a send or receive exit is called. This is useful for identifying which channel flows include user data and might require processing such as encryption or digital signing.

The following table shows the data that appears in byte 10 of the channel flow when an API call is being processed.

**Note:** These are not the only values of this byte. There are other **reserved** values.

Table 117. Identifying API calls

| API call                    | Value of byte 10 for request | Value of byte 10 for reply |
|-----------------------------|------------------------------|----------------------------|
| MQCONN <sup>1, 2</sup>      | X'81'                        | X'91'                      |
| MQDISC <sup>1</sup>         | X'82'                        | X'92'                      |
| MQOPEN <sup>3</sup>         | X'83'                        | X'93'                      |
| MQCLOSE                     | X'84'                        | X'94'                      |
| MQGET <sup>4</sup>          | X'85'                        | X'95'                      |
| MQPUT <sup>4</sup>          | X'86'                        | X'96'                      |
| MQPUT1 request <sup>4</sup> | X'87'                        | X'97'                      |
| MQSET request               | X'88'                        | X'98'                      |
| MQINQ request               | X'89'                        | X'99'                      |
| MQCMIT request              | X'8A'                        | X'9A'                      |
| MQBACK request              | X'8B'                        | X'9B'                      |
| MQSTAT request              | X'8D'                        | X'9D'                      |
| MQSUB request               | X'8E'                        | X'9E'                      |
| MQSUBRQ request             | X'8F'                        | X'9F'                      |
| xa_start request            | X'A1'                        | X'B1'                      |
| xa_end request              | X'A2'                        | X'B2'                      |
| xa_open request             | X'A3'                        | X'B3'                      |
| xa_close request            | X'A4'                        | X'B4'                      |
| xa_prepare request          | X'A5'                        | X'B5'                      |
| xa_commit request           | X'A6'                        | X'B6'                      |
| xa_rollback request         | X'A7'                        | X'B7'                      |
| xa_forget request           | X'A8'                        | X'B8'                      |
| xa_recover request          | X'A9'                        | X'B9'                      |
| xa_complete request         | X'AA'                        | X'BA'                      |

**Notes:**

1. The connection between the client and server is initiated by the client application using MQCONN. Therefore, for this command in particular, there are several other network flows. The same applies to MQDISC, which terminates the network connection.
2. MQCONNX is treated in the same way as MQCONN for the purposes of the client-server connection.
3. If a large distribution list is opened, there might be more than one network flow per MQOPEN call in order to pass all the required data to the SVRCONN MCA.
4. Large messages can exceed the transmission segment size. If this happens there can be many network flows resulting from a single API call.

## Connecting a client to a queue-sharing group

z/OS

You can connect a client to a queue-sharing group by creating an MQI channel between a client and a queue manager on a server that is a member of a queue-sharing group.

A queue-sharing group is formed by a set of queue-managers that can access the same set of shared queues. For more information about shared queues, see [Shared queues and queue-sharing groups](#).

A client putting to a shared queue can connect to any member of the queue-sharing group. The benefits of connecting to a queue-sharing group are possible increases in front-end and back-end availability, and increased capacity. You can connect to a specific queue manager or to the generic interface.

Connecting directly to a queue manager in a queue-sharing group gives the benefit that you can put messages to a shared target queue, which increases back-end availability.

Connecting to the generic interface of a queue-sharing group opens a session with one of the queue managers in the group. This increases front-end availability, because the client queue manager can connect with any queue-manager in the group. You connect to the group using the generic interface when you do not want to connect to a specific queue manager within the queue-sharing group.

The generic interface can be a Sysplex Distributor VIPA address or a VTAM generic resource name, or another common interface to the queue-sharing group. For more details on setting up a generic interface, see [Setting up communication for IBM MQ for z/OS using queue-sharing groups](#).

To connect to the generic interface of a queue-sharing group you need to create channel definitions that can be accessed by any queue manager in the group. To do this you need to have the same definitions on each queue manager in the group.

Define the SVRCONN channel as follows:

```
DEFINE CHANNEL(CHANNEL1) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
QSGDISP(GROUP)
```

Channel definitions on the server are stored in a shared Db2 repository. Each queue manager in the queue-sharing group makes a local copy of the definition, ensuring that you will always connect to the correct server-connection channel when you issue an MQCONN or MQCONNX call.

Define the CLNTCONN channel as follows:

```
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(VIPA address) QMNAME(QSG1) +
DESCR('Client-connection to Queue Sharing Group QSG1') QSGDISP(GROUP)
```

Because the generic interface of the queue-sharing group is stored in the CONNNAME field in the client-connection channel, you can now connect to any queue manager in the group, and put to shared queues owned by that group.

## Configuring a client using a configuration file

Configure your clients by using attributes in a text file. These attributes can be overridden by environment variables or in other platform-specific ways.

You can configure your IBM MQ MQI clients by using a text file, the IBM MQ MQI client configuration file. This file is similar to the queue manager configuration file, `qm.ini`, that is used on UNIX and Linux systems. The file contains a number of stanzas, each of which contains a number of lines of the format **attribute-name = value**.

The IBM MQ MQI client configuration file is generally named `mqclient.ini`, but you can choose to give it another name. Configuration information in this file applies to all platforms, and to clients that use the MQI, IBM MQ classes for Java, IBM MQ classes for JMS, IBM MQ classes for .NET, and XMS. Although the attributes in the IBM MQ MQI client configuration file apply to most IBM MQ clients, there are some attributes that are not read by managed .NET and XMS .NET clients, or by clients using either the IBM MQ classes for Java or the IBM MQ classes for JMS. For more information, see “Which IBM MQ clients can read each attribute” on page 729.

The configuration features apply to all connections a client application makes to any queue managers, rather than being specific to an individual connection to a queue manager. Attributes relating to a connection to an individual queue manager can be configured programmatically, for example by using an MQCD structure, or by using a Client Channel Definition Table (CCDT).

Environment variables that were supported in releases of IBM WebSphere MQ earlier than Version 7.0 continue to be supported, and where such an environment variable matches an equivalent value in the client configuration file, the environment variable overrides the client configuration file value.

For a client application that uses IBM MQ classes for JMS, you can also override the client configuration file in the following ways:

- By setting properties in the JMS configuration file
- By setting Java system properties, which also overrides the JMS configuration file.

For the .NET client, you can also override the client configuration file and the equivalent environment variables by using the .NET application configuration file.

You cannot set up multiple channel connections by using the client configuration file.

### Example client configuration file

```
Module Name: mqclient.ini
Type : IBM MQ MQI client configuration file
Function : Define the configuration of a client
##
##
Notes :
1) This file defines the configuration of a client
##
##
```

```
ClientExitPath:
 ExitsDefaultPath=/var/mqm/exits
 ExitsDefaultPath64=/var/mqm/exits64
```

```
TCP:
 Library1=DLLName1
 KeepAlive = Yes
 ClntSndBuffSize=32768
 ClntRcvBuffSize=32768
 Connect_Timeout=0
```

```
MessageBuffer:
 MaximumSize=-1
 Updatepercentage=-1
 PurgeTime=0
```

```
LU62:
 TPName
 Library1=DLLName1
 Library2=DLLName2
```

```
PreConnect:
 Module=amqldapi
 Function=myFunc
 Data=ldap://myLDAPServer.com:389/cn=wmq,ou=ibm,ou=com
 Sequence=1
```

```
CHANNELS:
 DefRecon=YES
 ServerConnectionParms=SALES.SVRCONN/TCP/hostname.x.com(1414)
```

### Related reference:

“Using IBM MQ environment variables” on page 747

This section describes the environment variables that you can use with IBM MQ MQI client applications.

“Changing queue manager configuration information” on page 782

The attributes described here modify the configuration of an individual queue manager. They override any settings for IBM MQ.

## Location of the client configuration file

An IBM MQ MQI client configuration file can be held in a number of locations.

A client application uses the following search path to locate the IBM MQ MQI client configuration file:

1. The location specified by the environment variable MQCLNTCF.

The format of this environment variable is a full URL. This means the file name might not necessarily be `mqclient.ini` and facilitates placing the file on a network attached file-system.

Note the following:

- C, .NET and XMS clients support only the `file:` protocol; the `file:` protocol is assumed if the URL string does not begin with `protocol:`
- To allow for Java 1.4.2 JREs, which do not support reading environment variables, the MQCLNTCF environment variable can be overridden with an MQCLNTCF Java System Property.

2. A file called `mqclient.ini` in the present working directory of the application.

3. A file called `mqclient.ini` in the IBM MQ data directory for Windows, UNIX and Linux systems.

Note the following:

- The IBM MQ data directory does not exist on certain platforms, for example, IBM i and z/OS, or in cases where the client has been supplied with another product.
- On UNIX and Linux systems, the directory is `/var/mqm`
- On Windows platforms you configure the environment variable `MQ_FILE_PATH` during installation, to point at the data directory. It is normally `C:\ProgramData\IBM\MQ`
- To allow for Java 1.4.2 JREs that do not support reading environment variables you can manually override the `MQ_FILE_PATH` environment variable with an `MQ_FILE_PATH` Java System Property.

4. A file called `mqclient.ini` in a standard directory appropriate to the platform, and accessible to users:

- For all Java clients this is the value of the `user.home` Java System Property.
- For C clients on UNIX and Linux platforms this is the value of the `HOME` environment variable.
- For C clients on Windows this is the concatenated values of the `HOMEDRIVE` and `HOMEPATH` environment variables.



**Note:** For the IBM MQ client for HP Integrity NonStop Server, the mqclient.ini file must be located in the OSS file system. Guardian applications must either place the mqclient.ini file in the IBM MQ data directory or set the MQCLNTCF environment variable to a location in the OSS file system.

### Which IBM MQ clients can read each attribute

Most of the attributes in the IBM MQ MQI client configuration file can be used by the C client, and the unmanaged .NET clients. However, there are some attributes that are not read by managed .NET and XMS .NET clients, or by clients using either the IBM MQ classes for Java or the IBM MQ classes for JMS.

Table 118. Which attributes apply to each type of client

| mqclient.ini stanza name and attributes | Description                                                                                                                                                                                   | C and unmanaged .NET | Java | JMS | Managed .NET and XMS .NET |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|------|-----|---------------------------|
| <b>CHANNELS stanza</b>                  |                                                                                                                                                                                               |                      |      |     |                           |
| CCSID                                   | The coded character set number to be used.                                                                                                                                                    | Yes                  | No   | No  | Yes                       |
| ChannelDefinitionDirectory              | The directory path to the file containing the client channel definition table.                                                                                                                | Yes                  | No   | No  | Yes                       |
| ChannelDefinitionFile                   | The name of the file containing the client channel definition table.                                                                                                                          | Yes                  | No   | No  | No                        |
| ReconDelay                              | An administrative option to configure reconnect delay for client programs that can auto-reconnect.                                                                                            | Yes                  | No   | Yes | No                        |
| DefRecon                                | An administrative option to enable client programs to automatically reconnect, or to disable the automatic reconnection of a client program that has been written to reconnect automatically. | Yes                  | No   | Yes | No                        |
| MQReconnectTime                         | The timeout in seconds to reconnect to a client.                                                                                                                                              | Yes                  | No   | No  | No                        |

Table 118. Which attributes apply to each type of client (continued)

| mqclient.ini stanza name and attributes | Description                                                                                                                                                                           | C and unmanaged .NET | Java | JMS | Managed .NET and XMS .NET |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|------|-----|---------------------------|
| ServerConnectionPath                    | Specifies the location of the IBM MQ server and the communication method to be used.                                                                                                  | Yes                  | No   | No  | Yes                       |
| Put1DefaultAlwaysSync                   | Controls the behavior of the MQPUT1 function call with the option MQPMO_RESPONSE_ASYNC.                                                                                               | Yes                  | Yes  | Yes | No                        |
| PasswordProtection                      | Allows you to set protected passwords in the MQCSP structure, rather than using SSL or TLS.                                                                                           | Yes                  | Yes  | Yes | No                        |
| <b>ClientExitPath stanza</b>            |                                                                                                                                                                                       |                      |      |     |                           |
| ExitsDefaultPath                        | Specifies the location of 32-bit channel exits for clients.                                                                                                                           | Yes                  | No   | No  | Yes                       |
| ExitsDefaultPath64                      | Specifies the location of 64-bit channel exits for clients.                                                                                                                           | Yes                  | No   | No  | Yes                       |
| JavaExitsClassPath                      | The values to be added to the classpath when a Java exit is run.                                                                                                                      | No                   | Yes  | Yes | No                        |
| <b>JMQI stanza</b>                      |                                                                                                                                                                                       |                      |      |     |                           |
| useMQCSPauthentication                  | Controls whether IBM MQ classes for Java and IBM MQ classes for JMS applications should use Compatibility mode or MQCSP authentication mode when authenticating with a queue manager. | No                   | Yes  | Yes | No                        |
| <b>MessageBuffer stanza</b>             |                                                                                                                                                                                       |                      |      |     |                           |

Table 118. Which attributes apply to each type of client (continued)

| <b>mqclient.ini stanza name and attributes</b> | <b>Description</b>                                                                                                                                                      | <b>C and unmanaged .NET</b> | <b>Java</b> | <b>JMS</b> | <b>Managed .NET and XMS .NET</b> |  |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|-------------|------------|----------------------------------|--|
| MaximumSize                                    | Size, in kilobytes, of the read-ahead buffer, in the range 1 through 999 999.                                                                                           | Yes                         | Yes         | Yes        | Yes                              |  |
| PurgeTime                                      | Interval, in seconds, after which messages left in the read-ahead buffer are purged.                                                                                    | Yes                         | Yes         | Yes        | Yes                              |  |
| UpdatePercentage                               | The update percentage value, in the range of 1 - 100, used in calculating the threshold value to determine when a client application makes a new request to the server. | Yes                         | Yes         | Yes        | Yes                              |  |
| <b>PreConnect stanza</b>                       |                                                                                                                                                                         |                             |             |            |                                  |  |
| Data                                           | URL of the repository where connection definitions are stored.                                                                                                          | Yes                         | No          | No         | No                               |  |
| Function                                       | Name of the functional entry point into the library that contains the PreConnect exit code.                                                                             | Yes                         | No          | No         | No                               |  |
| Module                                         | The name of the module containing the API exit code.                                                                                                                    | Yes                         | No          | No         | No                               |  |
| Sequence                                       | The sequence in which this exit is called relative to other exits.                                                                                                      | Yes                         | No          | No         | No                               |  |
| <b>SSL stanza</b>                              |                                                                                                                                                                         |                             |             |            |                                  |  |


Table 118. Which attributes apply to each type of client (continued)

| mqclient.ini stanza name and attributes | Description                                                                                                                                         | C and unmanaged .NET | Java | JMS | Managed .NET and XMS .NET |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|------|-----|---------------------------|
| CDPCheckExtensions                      | Specifies whether SSL or TLS channels on this queue manager try to check CDP servers that are named in CrlDistributionPoint certificate extensions. | Yes                  | No   | No  | No                        |
| CertificateLabel                        | The certificate label of the channel definition.                                                                                                    | Yes                  | No   | No  | No                        |
| CertificateValPolicy                    | Determines the type of certificate validation used.                                                                                                 | Yes                  | No   | No  | No                        |
| ClientRevocationCheck                   | Determines how certificate revocation checking is configured if the client connect call uses an SSL/TLS channel.                                    | Yes                  | No   | No  | No                        |
| EncryptionPolicySuiteB                  | Determines whether a channel uses Suite-B compliant cryptography and what level of strength is to be used.                                          | Yes                  | No   | No  | No                        |
| OCSPAuthentication                      | Defines the behavior of IBM MQ when OCSP is enabled and the OCSP revocation check is unable to determine the certificate revocation status.         | Yes                  | No   | No  | No                        |
| OCSPCheckExtensions                     | Controls whether IBM MQ acts on AuthorityInfoAccess certificate extensions.                                                                         | Yes                  | No   | No  | No                        |

Table 118. Which attributes apply to each type of client (continued)

| <b>mqclient.ini stanza name and attributes</b> | <b>Description</b>                                                                                                          | <b>C and unmanaged .NET</b> | <b>Java</b> | <b>JMS</b> | <b>Managed .NET and XMS .NET</b> |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|-----------------------------|-------------|------------|----------------------------------|
| SSLCryptoHardware                              | Sets the parameter string required to configure PKCS #11 cryptographic hardware present on the system.                      | Yes                         | No          | No         | No                               |
| SSLFipsRequired                                | Specifies whether only FIPS-certified algorithms are to be used if cryptography is carried out in IBM MQ.                   | Yes                         | No          | No         | No                               |
| SSLHTTPProxyName                               | The string is either the host name or network address of the HTTP Proxy server that is to be used by GSKit for OCSP checks. | Yes                         | No          | No         | No                               |
| SSLKeyRepository                               | The location of the key repository that holds the user's digital certificate, in stem format.                               | Yes                         | No          | No         | No                               |
| SSLKeyResetCount                               | The number of unencrypted bytes sent and received on an SSL or TLS channel before the secret key is renegotiated.           | Yes                         | No          | No         | No                               |
| <b>TCP stanza</b>                              |                                                                                                                             |                             |             |            |                                  |
| ClntRcvBuffSize                                | The size in bytes of the TCP/IP receive buffer used by the client end of a client-connection server-connection channel.     | Yes                         | Yes         | Yes        | Yes                              |

Table 118. Which attributes apply to each type of client (continued)

| mqclient.ini stanza name and attributes                                                    | Description                                                                                                          | C and unmanaged .NET | Java | JMS | Managed .NET and XMS .NET |
|--------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|----------------------|------|-----|---------------------------|
| ClntSndBufferSize                                                                          | The size in bytes of the TCP/IP send buffer used by the client end of a client-connection server-connection channel. | Yes                  | Yes  | Yes | Yes                       |
| Connect_Timeout                                                                            | The number of seconds before an attempt to connect the socket times out.                                             | Yes                  | Yes  | Yes | No                        |
| IPAddressVersion                                                                           | Specifies which IP protocol to use for a channel connection.                                                         | Yes                  | No   | No  | Yes                       |
| KeepAlive                                                                                  | Switches the KeepAlive function on or off.                                                                           | Yes                  | Yes  | Yes | Yes                       |
|  Library | On Windows only, the name of the TCP/IP sockets DLL.                                                                 | Yes                  | No   | No  | No                        |

For the IBM MQ client for HP Integrity NonStop Server, you can use the TMF and TmfGateway stanzas to communicate with the TMF/Gateway.

## CHANNELS stanza of the client configuration file

Use the CHANNELS stanza to specify information about client channels.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see Which IBM MQ attributes can be read by each client.

The following attributes can be included in the CHANNELS stanza.

### CCSID = *number*

The coded character set number to be used.


This attribute can be read by C, unmanaged .NET, and managed .NET clients.

The CCSID number is equivalent to the MQCCSID environment parameter.

### ChannelDefinitionDirectory = *path*

The directory path to the file containing the client channel definition table.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

 On Windows systems, the default is the IBM MQ data and log files directory, typically C:\ProgramData\IBM\MQ.

  On UNIX and Linux systems, the default is /var/mqm.

The ChannelDefinitionDirectory path is equivalent to the MQCHLLIB environment parameter.

**ChannelDefinitionFile** = *filename* | AMQCLCHL.TAB

The name of the file containing the client channel definition table.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

The client channel definition table is equivalent to the MQCHLTAB environment parameter.

**ReconDelay** = (*delay[,rand]*) (*delay[,rand]*)...

The ReconDelay attribute provides an administrative option to configure reconnect delay for client programs that can auto-reconnect.

This attribute can be read by C, unmanaged .NET, and IBM MQ classes for JMS clients.

Here is an example configuration:

```
ReconDelay=(1000,200)(2000,200)(4000,1000)
```

The example shown defines an initial delay of one second, plus a random interval of up to 200 milliseconds. The next delay is two seconds plus a random interval of up to 200 milliseconds. All subsequent delays are four seconds, plus a random interval of up to 1000 milliseconds.

**DefRecon** = *NO* | *YES* | *QMGR* | *DISABLED*

The DefRecon attribute provides an administrative option to enable client programs to automatically reconnect, or to disable the automatic reconnection of a client program that has been written to reconnect automatically. You might opt to set the latter if a program uses an option, such as MQPMO\_LOGICAL\_ORDER, that is incompatible with reconnection.

This attribute can be read by C, unmanaged .NET, and IBM MQ classes for JMS clients.

Automatic client reconnection is not supported by IBM MQ classes for Java.

The interpretation of the DefRecon options depends on whether an MQCNO\_RECONNECT\_\* value is also set in the client program, and what value is set.

If the client program connects using MQCONN, or sets the MQCNO\_RECONNECT\_AS\_DEF option using MQCONNX, the reconnect value set by DefRecon takes effect. If no reconnect value is set in the program, or by the DefRecon option, the client program is not reconnected automatically.

**NO** Unless overridden by MQCONNX, the client is not reconnected automatically.

**YES** Unless overridden by MQCONNX, the client reconnects automatically.

**QMGR** Unless overridden by MQCONNX, the client reconnects automatically, but only to the same queue manager. The QMGR option has the same effect as MQCNO\_RECONNECT\_Q\_MGR.

**DISABLED**

Reconnection is disabled, even if requested by the client program using the MQCONNX MQI call.

The automatic client reconnection depends on two values:

- The reconnect option set in the application
- DefRecon value in the mqclient.ini file

Table 119. Automatic reconnection depends on the values set in the application and in the mqclient.ini file

| DefRecon value in the mqclient.ini | Reconnection options set in the application |                       |                        |                          |
|------------------------------------|---------------------------------------------|-----------------------|------------------------|--------------------------|
|                                    | MQCNO_RECONNECT                             | MQCNO_RECONNECT_Q_MGR | MQCNO_RECONNECT_AS_DEF | MQCNO_RECONNECT_DISABLED |
| NO                                 | YES                                         | QMGR                  | NO                     | NO                       |
| YES                                | YES                                         | QMGR                  | YES                    | NO                       |
| QMGR                               | YES                                         | QMGR                  | QMGR                   | NO                       |
| DISABLED                           | NO                                          | NO                    | NO                     | NO                       |

### MQReconnectTimeout

The timeout in seconds to reconnect to a client. The default value is 1800 seconds (30 minutes).

This attribute can be read by C and unmanaged .NET clients.

IBM MQ classes for JMS clients can specify a timeout to reconnect using the connection factory property CLIENTRECONNECTTIMEOUT. The default value for this property is 1800 seconds (30 minutes).

### ServerConnectionParms

ServerConnectionParms is equivalent to the MQSERVER environment parameter and specifies the location of the IBM MQ server and the communication method to be used.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

The ServerConnectionParms attribute defines only a simple channel. You cannot use it to define an SSL channel or a channel with channel exits. It is a string of the format *ChannelName/TransportType/ConnectionName*, *ConnectionName* must be a fully qualified network name. *ChannelName* cannot contain the forward slash (/) character because this character is used to separate the channel name, transport type, and connection name.

When ServerConnectionParms is used to define a client channel, a maximum message length of 100 MB is used. Therefore the maximum message size in effect for the channel is the value specified in the SVRCONN channel on the server.

Note that only a single client channel connection can be made. For example, if you have two entries:

```
ServerConnectionParms=R1.SVRCONN/TCP/localhost(1963)
ServerConnectionParms=R2.SVRCONN/TCP/localhost(1863)
```

only the second one is used.

Specify *ConnectionName* as a comma-separated list of names for the stated transport type. Generally, only one name is required. You can provide multiple *hostnames* to configure multiple connections with the same properties. The connections are tried in the order that they are specified in the connection list until a connection is successfully established. If no connection is successful, the client starts to process again. Connection lists are an alternative to queue manager groups to configure connections for reconnectable clients.

### Put1DefaultAlwaysSync = NO | YES

Controls the behavior of the MQPUT1 function call with the option MQPMO\_RESPONSE\_AS\_Q\_DEF.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, and IBM MQ classes for JMS clients.

**NO** If MQPUT1 is set with MQPMO\_SYNCPOINT, it behaves as MQPMO\_ASYNC\_RESPONSE. Similarly, if MQPUT1 is set with MQPMO\_NO\_SYNCPOINT, it behaves as MQPMO\_SYNC\_RESPONSE. This is the default value.

**YES** MQPUT1 behaves as if MQPMO\_SYNC\_RESPONSE is set, regardless of whether MQPMO\_SYNCPOINT or MQPMO\_NO\_SYNCPOINT is set.



**PasswordProtection** = *Compatible | always | optional*

From IBM MQ Version 8.0, allows you to set protected passwords in the MQCSP structure, rather than using SSL or TLS.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, and IBM MQ classes for JMS clients.

MQCSP password protection is useful for test and development purposes as using MQCSP password protection is simpler than setting up SSL/TLS encryption, but not as secure. For more information, see MQCSP password protection.

**Related information:**

Connecting IBM MQ MQI applications to queue managers

## **ClientExitPath stanza of the client configuration file**

Use the ClientExitPath stanza to specify the default locations of channel exits on the client.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see Which IBM MQ attributes can be read by each client.

The following attributes can be included in the ClientExitPath stanza.

**ExitsDefaultPath** = *string*

Specifies the location of 32-bit channel exits for clients.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

**ExitsDefaultPath64** = *string*

Specifies the location of 64-bit channel exits for clients.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

**JavaExitsClassPath** = *string*

The values to be added to the classpath when a Java exit is run. This is ignored by exits in any other language.

This attribute can be read by IBM MQ classes for Java and IBM MQ classes for JMS clients.

In the JMS configuration file, the JavaExitsClassPath name is given the standard `com.ibm.mq.cfg.` prefix and this full name is also used on the IBM WebSphere MQ Version 7.0 or later system property. At Version 6.0 this attribute was specified using system property `com.ibm.mq.exitClasspath`, which was documented in the Version 6.0 readme. The use of `com.ibm.mq.exitClasspath` is deprecated. If both `JavaExitsClassPath` and `exitClasspath` are present, `JavaExitsClassPath` is honored. If only `exitClasspath` usage is present, it is still honored in IBM WebSphere MQ Version 7.0 or later.

## JMQI stanza of the client configuration file

Use the JMQI stanza to specify configuration parameters for the Java Message Queuing Interface (JMQUI) used by the IBM MQ classes for Java and IBM MQ classes for JMS.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see Which IBM MQ attributes can be read by each client.

The following attribute can be included in the JMQI stanza:

**useMQCSPauthentication = NO | YES**

Controls whether IBM MQ classes for Java and IBM MQ classes for JMS applications should use Compatibility mode or MQCSP authentication mode when authenticating with a queue manager.

This attribute can be read by IBM MQ classes for Java, and IBM MQ classes for JMS clients.

This attribute can have the following values:

**NO** Use compatibility mode when authenticating with a queue manager. This is the default value.

**YES** Use MQCSP authentication mode when authenticating with a queue manager.

For more information about Compatibility mode and MQCSP authentication mode, see Connection authentication with the Java client.

## LU62, NETBIOS, and SPX stanzas of the client configuration file

Windows

On Windows systems only, use these stanzas to specify configuration parameters for the specified network protocols.

### LU62 stanza

Use the LU62 stanza to specify SNA LU 6.2 protocol configuration parameters. The following attributes can be included in this stanza:

**Library1 = DLLName | WCPIC32**  
The name of the APPC DLL.

**Library2 = DLLName | WCPIC32**  
The same as Library1, used if the code is stored in two separate libraries. .

**TPName**  
The TP name to start on the remote site.

### NETBIOS stanza

Use the NETBIOS stanza to specify NetBIOS protocol configuration parameters. The following attributes can be included in this stanza:

**AdapterNum = number | 0**  
The number of the LAN adapter.

**Library1 = DLLName | NETAPI32**  
The name of the NetBIOS DLL.

**LocalName = name**  
The name by which this computer is known on the LAN.  
This is equivalent to the MQNAME environment parameter.

**NumCmds** = *number* | 1

How many commands to allocate.

**NumSess** = *number* | 1

How many sessions to allocate.

## SPX stanza

Use the SPX stanza to specify SPX protocol configuration parameters. The following attributes can be included in this stanza:

**BoardNum** = *number* | 0

The LAN adapter number.

**KeepAlive** = YES | NO

Switch the KeepAlive function on or off.

KeepAlive = YES causes SPX to check periodically that the other end of the connection is still available. If it is not, the channel is closed.

**Library1** = *DLLName* | WSOCK32.DLL

The name of the SPX DLL.

**Library2** = *DLLName* | WSOCK32.DLL

The same as Library1, used if the code is stored in two separate libraries.

**Socket** = *number* | 5E86

The SPX socket number in hexadecimal notation.

## MessageBuffer stanza of the client configuration file

Use the MessageBuffer stanza to specify information about message buffers.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see Which IBM MQ attributes can be read by each client.

The following attributes can be included in the MessageBuffer stanza:

**MaximumSize** = *integer* | 1

Size, in kilobytes, of the read-ahead buffer, in the range 1 - 999 999.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

The following special values exist:

-1 The client determines the appropriate value.

0 Read ahead is disabled for the client.

**PurgeTime** = *integer* | 600

Interval, in seconds, after which messages left in the read-ahead buffer are purged.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

If the client application is selecting messages based on MsgId or CorrelId, it is possible that the read ahead buffer might contain messages that are sent to the client with a previously requested MsgId or CorrelId. These messages would then be stranded in the read ahead buffer until an MQGET is issued with an appropriate MsgId or CorrelId. You can purge messages from the read ahead buffer by setting PurgeTime. Any messages that have remained in the read ahead buffer for longer than the purge interval are automatically purged. These messages have already been removed from the queue on the queue manager, so unless they are being browsed, they are lost.

The valid range is in the range 1 - 999 999 seconds, or the special value 0, meaning that no purge takes place.

**UpdatePercentage** = *integer* | -1

The update percentage value, in the range of 1 - 100, used in calculating the threshold value to determine when a client application makes a new request to the server. The special value -1 indicates that the client determines the appropriate value.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

The client periodically sends a request to the server indicating how much data the client application has consumed. A request is sent when the number of bytes, *n*, retrieved by the client by way of MQGET calls exceeds a threshold *T*. *n* is reset to zero each time that a new request is sent to the server.

The threshold *T* is calculated as follows:

$$T = \text{Upper} - \text{Lower}$$

Upper is the same as the read-ahead buffer size, specified by the MaximumSize attribute, in Kilobytes. Its default is 100 Kb.

Lower is lower than Upper, and is specified by the UpdatePercentage attribute. This attribute is a number in the range 1 - 100, and has a default of 20. Lower is calculated as follows:

$$\text{Lower} = \text{Upper} \times \text{UpdatePercentage} / 100$$

**Example 1:**

The MaximumSize and UpdatePercentage attributes take their defaults of 100 Kb and 20 Kb.

The client calls MQGET to retrieve a message, and does so repeatedly. This continues until MQGET has consumed *n* bytes.

Using the calculation

$$T = \text{Upper} - \text{Lower}$$

*T* is (100 - 20) = 80 Kb.

So when MQGET calls have removed 80 Kb from a queue, the client makes a new request automatically.

**Example 2:**

The MaximumSize attribute takes its default of 100 Kb, and a value of 40 is chosen for UpdatePercentage.

The client calls MQGET to retrieve a message, and does so repeatedly. This continues until MQGET has consumed *n* bytes.

Using the calculation

$$T = \text{Upper} - \text{Lower}$$

*T* is (100 - 40) = 60 Kb

So when MQGET calls have removed 60 Kb from a queue, the client makes a new request automatically. This is sooner than in EXAMPLE 1 where the defaults were used.

Therefore, choosing a larger threshold *T* tends to decrease the frequency at which requests are sent from client to server. Conversely, choosing a smaller threshold *T* tends to increase the frequency of requests that are sent from client to server.

However, choosing a large threshold *T* can mean that the performance gain of read ahead is reduced as the chance of the read ahead buffer becoming empty can increase. When this happens an MQGET call might have to pause, waiting for data to arrive from the server.

## PreConnect stanza of the client configuration file

Use the PreConnect stanza to configure the PreConnect exit in the `mqclient.ini` file.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see *Which IBM MQ attributes can be read by each client*.

The following attributes can be included in the PreConnect stanza:

**Data** = *<URL>*

URL of the repository where connection definitions are stored.

This attribute can be read by C and unmanaged .NET clients.

For example, when using an LDAP server:

```
Data = ldap://myLDAPServer.com:389/cn=wmq,ou=ibm,ou=com
```

**Function** = *<myFunc>*

Name of the functional entry point into the library that contains the PreConnect exit code.

This attribute can be read by C and unmanaged .NET clients.

The function definition adheres to the PreConnect exit prototype `MQ_PRECONNECT_EXIT`.

The maximum length of this field is `MQ_EXIT_NAME_LENGTH`.

**Module** = *<amqldapi>*

The name of the module containing the API exit code.

This attribute can be read by C and unmanaged .NET clients.

If this field contains the full path name of the module, it is used as is.

**Sequence** = *<sequence\_number>*

The sequence in which this exit is called relative to other exits. An exit with a low sequence number is called before an exit with a higher sequence number. There is no need for the sequence numbering of exits to be continuous; a sequence of 1, 2, 3 has the same result as a sequence of 7, 42, 1096. This attribute is an unsigned numeric value.

This attribute can be read by C and unmanaged .NET clients.

Multiple PreConnect stanzas can be defined within the `mqclient.ini` file. The processing order of each exit is determined by the Sequence attribute of the stanza.

**Related information:**

Referencing connection definitions using a pre-connect exit from a repository

## SSL stanza of the client configuration file

Use the SSL stanza to specify information about the use of SSL or TLS.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see *Which IBM MQ attributes can be read by each client*.

The following attributes can be included in the SSL stanza:

**CDPCheckExtensions**= YES | NO

CDPCheckExtensions specifies whether SSL or TLS channels on this queue manager try to check CDP servers that are named in `CrlDistributionPoint` certificate extensions.

This attribute can be read by C and unmanaged .NET clients.

This attribute has the following possible values:

- YES: SSL or TLS channels try to check CDP servers to determine whether a digital certificate is revoked.
- NO: SSL or TLS channels do not try to check CDP servers. This value is the default.

**CertificateLabel** = *string*

The certificate label of the channel definition.

This attribute can be read by C and unmanaged .NET clients.

For more information, see Certificate label (CERTLABEL).

**CertificateValPolicy** = *string*

Determines the type of certificate validation used.

This attribute can be read by C and unmanaged .NET clients.

This attribute has the following possible values:

**ANY** Use any certificate validation policy supported by the underlying secure sockets library. This setting is the default setting.

**RFC5280**

Use only certificate validation which complies with the RFC 5280 standard.

**ClientRevocationChecks** = **REQUIRED** | **OPTIONAL** | **DISABLED**


Determines how certificate revocation checking is configured if the client connect call uses an SSL/TLS channel. See also **OCSPAuthentication**.

This attribute can be read by C and unmanaged .NET clients.

This attribute has the following possible values:

**REQUIRED (default)**

Attempts to load certificate revocation configuration from the CCDT and perform revocation checking as configured. If the CCDT file cannot be opened or it is not possible to validate the certificate (because an OCSP or CRL server is not available, for example) the MQCONN call fails. No revocation checking is performed if the CCDT contains no revocation configuration but this does not cause the channel to fail.

 **Windows** On Windows systems, you can also use Active Directory for CRL revocation checking. You cannot use Active Directory for OCSP revocation checking.

**OPTIONAL**

As for REQUIRED, but if it is not possible to load the certificate revocation configuration, the channel does not fail.

**DISABLED**

No attempt is made to load certificate revocation configuration from the CCDT and no certificate revocation checking is done.

**Note:** If you are using MQCONNX rather than MQCONN calls, you might choose to supply authentication information records (MQAIR) via the MQSCO. The default behavior with MQCONNX is therefore not to fail if the CCDT file cannot be opened but to assume that you are supplying an MQAIR (even if you choose not to do so).

**EncryptionPolicySuiteB** = *string*

Determines whether a channel uses Suite-B compliant cryptography and what level of strength is to be used.

This attribute can be read by C and unmanaged .NET clients.

This attribute has the following possible values:

**NONE**

Suite-B compliant cryptography is not used. This setting is the default setting.

**128\_BIT,192\_BIT**

Sets the security strength to both 128-bit and 192-bit levels.

**128\_BIT**

Sets the security strength to 128-bit level.

**192\_BIT**

Sets the security strength to 192-bit level.

**OCSPAuthentication = OPTIONAL | REQUIRED | WARN**

Defines the behavior of IBM MQ when OCSP is enabled and the OCSP revocation check is unable to determine the certificate revocation status. See also **ClientRevocationChecks**.

This attribute can be read by C and unmanaged .NET clients.

This attribute has the following possible values:

**OPTIONAL**

Any certificate with a revocation status that cannot be determined by OCSP checking is accepted and no warning or error message is generated. The SSL or TLS connection continues as if no revocation check had been made.

**REQUIRED**

OCSP checking must yield a definitive revocation result for every SSL or TLS certificate which is checked. Any SSL or TLS certificate with a revocation status that cannot be verified is rejected with an error message. If queue manager SSL event messages are enabled, an MQRC\_CHANNEL\_SSL\_ERROR message with a ReasonQualifier of MQRQ\_SSL\_HANDSHAKE\_ERROR is generated. The connection is closed.

This value is the default value.

**WARN**

A warning is reported in the queue manager error logs if an OCSP revocation check is unable to determine the revocation status of any SSL or TLS certificate. If queue manager SSL event messages are enabled, an MQRC\_CHANNEL\_SSL\_WARNING message with a ReasonQualifier of MQRQ\_SSL\_UNKNOWN\_REVOCATION is generated. The connection is allowed to continue.

**OCSPCheckExtensions = YES | NO**

Controls whether IBM MQ acts on AuthorityInfoAccess certificate extensions.

This attribute can be read by C and unmanaged .NET clients.

If the value is set to NO, IBM MQ ignores AuthorityInfoAccess certificate extensions and does not attempt an OCSP security check. The default value is YES.

**SSLCryptoHardware = *string***

Sets the parameter string required to configure PKCS #11 cryptographic hardware present on the system.

This attribute can be read by C and unmanaged .NET clients.

Specify a string in the following format: **GSK\_PKCS11** = *driver path and filename ; token label ; token password ; symmetric cipher setting ;*

For example: **GSK\_PKCS11**=/usr/lib/pkcs11/PKCS11\_API.so;tokenlabel;passw0rd;SYMMETRIC\_CIPHER\_ON

The driver path is an absolute path to the shared library providing support for the PKCS #11 card. The driver file name is the name of the shared library. An example of the value required for the PKCS #11 driver path and file name is /usr/lib/pkcs11/PKCS11\_API.so. To access symmetric cipher operations through GSKit, specify the symmetric cipher setting parameter. This parameter can have either of the following values:

**SYMMETRIC\_CIPHER\_OFF**

Do not access symmetric cipher operations. This setting is the default setting.

## **SYMMETRIC\_CIPHER\_ON**

Access symmetric cipher operations.

The maximum length of the string is 256 characters. The default value is blank. If you specify a string that is not in the correct format, an error is generated.

## **SSLFipsRequired = YES | NO**

Specifies whether only FIPS-certified algorithms are to be used if cryptography is carried out in IBM MQ.

This attribute can be read by C, and unmanaged .NET clients.

If cryptographic hardware is configured, the cryptographic modules used are those modules provided by the hardware product. These might, or might not, be FIPS-certified to a particular level, depending on the hardware product in use.

## **SSLHTTPProxyName = string**

The string is either the host name or network address of the HTTP Proxy server that is to be used by GSKit for OCSP checks. This address can be followed by an optional port number, enclosed in parentheses. If you do not specify the port number, the default HTTP port, 80, is used.

This attribute can be read by C, and unmanaged .NET clients.

On the HP-UX PA-RISC and Sun Solaris SPARC platforms, and for 32-bit clients on AIX, the network address can be only an IPv4 address; on other platforms it can be an IPv4 or IPv6 address.

This attribute might be necessary if, for example, a firewall prevents access to the URL of the OCSP responder.

## **SSLKeyRepository = pathname**

The location of the key repository that holds the user's digital certificate, in stem format. That is, it includes the full path and the file name without an extension.

This attribute can be read by C, and unmanaged .NET clients.

## **SSLKeyResetCount = integer | 0**

The number of unencrypted bytes sent and received on an SSL or TLS channel before the secret key is renegotiated.

This attribute can be read by C, and unmanaged .NET clients.

The value must be in the range 0 - 999999999.

The default is 0, which means that secret keys are never renegotiated.

If you specify a value of 1 - 32768, SSL or TLS channels use a secret key reset count of 32768 (32Kb). This is to avoid excessive key resets, which would occur for small secret key reset values.



## TCP stanza of the client configuration file

Use the TCP stanza to specify TCP network protocol configuration parameters.

**Note:** The description of each attribute of this stanza indicates which IBM MQ clients can read that attribute. For a summary table for all IBM MQ MQI client configuration file stanzas, see Which IBM MQ attributes can be read by each client.

The following attributes can be included in the TCP stanza:

**ClntRcvBuffSize** = *number* | **0**

The size in bytes of the TCP/IP receive buffer used by the client end of a client-connection server-connection channel.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

A value of zero indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by IBM MQ. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**ClntSndBuffSize** = *number* | **0**

The size in bytes of the TCP/IP send buffer used by the client end of a client-connection server-connection channel.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

A value of zero indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by IBM MQ. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**Connect\_Timeout** = *number*

The number of seconds before an attempt to connect the socket times out. The default value of zero specifies that there is no connect timeout.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, and IBM MQ classes for JMS clients.

IBM MQ channel processes connect over nonblocking sockets. Therefore, if the other end of the socket is not ready, connect() returns immediately with *EINPROGRESS* or *EWOULDBLOCK*. Following this, connect will be attempted again, up to a total of 20 such attempts, when a communications error is reported.

If Connect\_Timeout is set to a non-zero value, IBM MQ waits for the stipulated period over select() call for the socket to get ready. This increases the chances of success of a subsequent connect() call. This option might be beneficial in situations where connects would require some waiting period, due to high load on the network.

There is no relationship between the Connect\_Timeout, ClntSndBuffSize, and ClntRcvBuffSize parameters.

**IPAddressVersion** = MQIPADDR\_IPV4 | MQIPADDR\_IPV6

Specifies which IP protocol to use for a channel connection.

This attribute can be read by C, unmanaged .NET, and managed .NET clients.

It has the possible string values of MQIPADDR\_IPV4 or MQIPADDR\_IPV6. These values have the same meanings as IPV4 and IPV6 in **ALTER QMGR IPADDRV**.

**KeepAlive** = **YES** | **NO**

Switch the KeepAlive function on or off. KeepAlive=YES causes TCP/IP to check periodically that the other end of the connection is still available. If it is not, the channel is closed.

This attribute can be read by C, unmanaged .NET, IBM MQ classes for Java, IBM MQ classes for JMS, and managed .NET clients.

**Windows** **Library1** = *DLLName* | **WSOCK32**  
( Windows only) The name of the TCP/IP sockets DLL.

This attribute can be read by C and unmanaged .NET clients.

## TMF and TmfGateway stanzas

Use the TMF and TMF/Gateway stanzas to specify the required configuration parameters for the IBM MQ client for HP Integrity NonStop Server to communicate with the TMF/Gateway.

If you want to use the HP NonStop Transaction Management Facility (TMF), then you must define a TMF stanza and one TmfGateway stanza for each queue manager with which you are communicating. All values are derived from your configuration.

The IBM MQ provided TMF/Gateway runs in a Pathway environment.

### TMF stanza

**PathMon** = *name*

The name of your defined Pathmon process that defines the server classes for the TMF/Gateway.

### TmfGateway stanza

The following attributes can be included in this stanza:

**QManager** = *name*

The name of the queue manager.

**Server** = *name*

The server class name for the TMF/Gateway configured for that queue manager.

## Example

Here is an example of a TMF stanza that is defined with two TmfGateway stanzas for two different queue managers on different servers:

```
TMF:
PathMon=$PSD1P
```

```
TmfGateway:
QManager=MQ5B
Server=MQ-MQ5B
```

```
TmfGateway:
QManager=MQ5C
Server=MQ-MQ5C
```

### Related concepts:

“Gateway process overview” on page 1187

The HP NonStop Transaction Management Facility (TMF) provides services to enable a gateway process to register as a resource manager. The IBM MQ for HP Integrity NonStop Server provided TMF/Gateway process runs under Pathway.

“Configuring the client initialization file” on page 1189

If you are using the HP NonStop Transaction Management Facility (TMF), you must have an IBM MQ client initialization file to enable your IBM MQ client for the HP Integrity NonStop Server to reach the TMF Gateway.

## Using IBM MQ environment variables


This section describes the environment variables that you can use with IBM MQ MQI client applications.

You can use environment variables in the following ways:

- Set the variables in your system profile to make a permanent change
- Issue a command from the command line to make a change for this session only
- To give one or more variables a particular value dependent on the application that is running, add commands to a command script file used by the application

IBM MQ uses default values for those variables that you have not set.

Commands are available on all the IBM MQ MQI client platforms unless otherwise stated.

**Note:**  IBM MQ for z/OS does not support any IBM MQ environment variables. If you are using this platform as your server, see Client channel definition table for information about how the client channel definition table is generated on z/OS. You can still use the IBM MQ environment variables on your client platform.

For each environment variable, use the command relevant to your platform to display the current setting or to reset the value of a variable.

On Windows, use the following commands:

- To remove the value of an environment variable, use the command `SET MQSERVER=`
- To display the current setting of an environment variable, use the command `SET MQSERVER`
- To display all environment variables for the session, use the command `set`

On UNIX and Linux systems, use the following commands:

- To remove the value of an environment variable, use the command `unset MQSERVER`
- To display the current setting of an environment variable, use the command `echo $MQSERVER`
- To display all environment variables for the session, use the command `set`

For information about the individual variables, see the following subtopics:

**Related concepts:**

“Configuring a client using a configuration file” on page 727

Configure your clients by using attributes in a text file. These attributes can be overridden by environment variables or in other platform-specific ways.

**Related information:**


Environment variables

**MQCCSID**

MQCCSID specifies the coded character set number to be used and overrides the CCSID value with which the server has been configured.

See Choosing client or server coded character set identifier (CCSID) for more information.

To set this variable use one of these commands:

- For Windows:  
SET MQCCSID=number
- For UNIX and Linux systems:  
export MQCCSID=number
-  For IBM i:  
ADDENVVAR ENVVAR(MQCCSID) VALUE(number)

**MQCERTLABL**

MQCERTLABL specifies the certificate label of the channel definition.

See Certificate label (CERTLABL) for more information.

**MQCERTVPOL**

MQCERTVPOL specifies the certificate validation policy used.

For more information about certificate validation policies in IBM MQ, see Certificate validation policies in IBM MQ.


This environment variable overrides the *CertificateValPolicy* setting in the SSL stanza of the client ini file. The variable can be set to one of two values:

**ANY** Use any certificate validation policy supported by the underlying secure sockets library.

**RFC5280**

Use only certificate validation which complies with the RFC 5280 standard.





To set this variable, use one of these commands:

- For Windows:  
SET MQCERTVPOL= *value*
- For UNIX and Linux systems:  
export MQCERTVPOL= *value*
-  For IBM i:  
ADDENVVAR ENVVAR(MQCERTVPOL) VALUE(*value*)





## MQCHLLIB

MQCHLLIB specifies the directory path to the file containing the client channel definition table (CCDT). The file is created on the server, but can be copied across to the IBM MQ MQI client workstation.





If MQCHLLIB is not set, the path for the client defaults to:

-  For Windows: `MQ_INSTALLATION_PATH`
-   For UNIX and Linux systems: `/var/mqm/`
-  For IBM i: `/QIBM/UserData/mqm/`



For the `crtmqm` and `strmqm` commands, the path defaults to one of two sets of paths. If `datapath` is set, the path defaults to one of the first set. If `datapath` is not set, the path defaults to one of the second set.

-  For Windows: `datapath\@ipcc`
-   For UNIX and Linux systems: `datapath/@ipcc`
-  For IBM i: `datapath/&ipcc`


Or:

-  For Windows: `MQ_INSTALLATION_PATH\data\qmgrs\qmgrname\@ipcc`
-   For UNIX and Linux systems: `/prefix/qmgrs/qmgrname/@ipcc`
-  For IBM i: `/prefix/qmgrs/qmgrname/&ipcc`

where:

- `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.
- If present, `datapath` is the value of DataPath defined in the queue manager stanza.
- `prefix` is the value of Prefix defined in the queue manager stanza. Prefix is typically `/var/mqm` on UNIX and Linux platforms , and `/QIBM/UserData/mqm/` on IBM i.
- `qmgrname` is the value of the Directory attribute defined in the queue manager stanza. The value might be different from the actual queue manager name. The value might have been altered to replace special characters.
- The queue manager stanza is defined in the `mqs.ini` file on  IBM, UNIX, and Linux, and in the registry on Windows

### Notes:

1.  If you are using IBM MQ for z/OS as your server, the file must be kept on the IBM MQ client workstation.
2. If set, MQCHLLIB overrides the path used to locate the CCDT.
3. Environment variables, such as MQCHLLIB, can be scoped to a process, or a job, or system-wide, in a platform-specific way.
4. If you set MQCHLLIB system-wide on a server, it sets the same path to the CCDT file for all the queue managers on the server. If you do not set the MQCHLLIB environment variable, the path is different for each queue manager. Queue managers read the value of MQCHLLIB, if it is set, on either the `crtmqm` or `strmqm` command.
5. If you create multiple queue managers on one server, the distinction is important, for the following reason. If you set MQCHLLIB system-wide, each queue manager updates the same CCDT file. The file contains the client-connection definitions from all the queue managers on the server. If the same definition exists on multiple queue managers, SYSTEM.DEF.CLNTCONN for example, the file contains the latest definition. When you create a queue manager, if MQCHLLIB is set, SYSTEM.DEF.CLNTCONN is updated in the CCDT. The update overwrites the SYSTEM.DEF.CLNTCONN created by a different queue manager. If you modified the earlier definition, your modifications are lost. For this reason, you must consider finding alternatives to setting MQCHLLIB as a system-wide environment variable on the server.

6. The MQSC and PCF NOREPLACE option on a client-connection definition does not check the contents of the CCDT file. A client-connection channel definition of the same name that was previously created, but not by this queue manager, is replaced, regardless of the NOREPLACE option. If the definition was previously created by the same queue manager, the definition is not replaced.
7. The command, **rcrmqobj -t clchltab** deletes and recreates the CCDT file. The file is recreated with only the client-connection definitions created on the queue manager that the command is running against.
8. Other commands that update the CCDT modify only the client-connection channels that have the same channel name. Other client-connection channels in the file are not altered.
9. The path for MQCHLLIB does not need quotations marks.




## Examples

To set this variable use one of these commands:

-  For Windows:  
SET MQCHLLIB=pathname

For example:

```
SET MQCHLLIB=C:\wmqtest
```


-   For UNIX and Linux systems:  
export MQCHLLIB=pathname
-  For IBM i:  
ADDENVVAR ENVVAR(MQCHLLIB) VALUE(pathname)

## MQCHLTAB

MQCHLTAB specifies the name of the file containing the client channel definition table (ccdt). The default file name is AMQCLCHL.TAB.

For information about where the client channel definition table is located on a server, see “Client channel definition table” on page 716.

To set this variable use one of these commands:

- On Windows:  
SET MQCHLTAB=filename
- On UNIX and Linux systems:  
export MQCHLTAB=filename
-  On IBM i:  
ADDENVVAR ENVVAR(MQCHLTAB) VALUE(filename)

For example:


```
SET MQCHLTAB=ccdf1.tab
```

In the same way as for the client, the MQCHLTAB environment variable on the server specifies the name of the client channel definition table.

## MQIPADDRV

MQIPADDRV specifies which IP protocol to use for a channel connection. It has the possible string values of "MQIPADDR\_IPv4" or "MQIPADDR\_IPv6". These values have the same meanings as IPv4 and IPv6 in ALTER QMGR IPADDRV. If it is not set, "MQIPADDR\_IPv4" is assumed.

To set this variable use one of these commands:

- For Windows:  
`SET MQIPADDRV=MQIPADDR_IPv4|MQIPADDR_IPv6`
- For UNIX and Linux systems:  
`export MQIPADDRV=MQIPADDR_IPv4|MQIPADDR_IPv6`
-  For IBM i:  
`ADDENVVAR ENVVAR(MQIPADDRV) VALUE(MQIPADDR_IPv4|MQIPADDR_IPv6)`

## MQNAME

MQNAME specifies the local NetBIOS name that the IBM MQ processes can use.

See “Defining a NetBIOS connection on Windows” on page 870 for a full description and for the rules of precedence on the client and the server.

To set this variable use this command:

```
SET MQNAME=Your_env_Name
```

For example:

```
SET MQNAME=CLIENT1
```

The NetBIOS on some platforms requires a different name (set by MQNAME) for each application if you are running multiple IBM MQ applications simultaneously on the IBM MQ MQI client.



## MQSERVER

MQSERVER environment variable is used to define a minimal channel. MQSERVER specifies the location of the IBM MQ server and the communication method to be used.

You cannot use MQSERVER to define an SSL channel or a channel with channel exits. For details of how to define an SSL channel, see Protecting channels with SSL.

*ConnectionName* must be a fully-qualified network name. The *ChannelName* cannot contain the forward slash (/) character because this character is used to separate the channel name, transport type, and connection name. When the MQSERVER environment variable is used to define a client channel, a maximum message length (MAXMSGL) of 100 MB is used. Therefore the maximum message size in effect for the channel is the value specified in the SVRCONN channel at the server.

To set this variable use one of these commands:

- For Windows:  
`SET MQSERVER=ChannelName/TransportType/ConnectionName`
- For UNIX and Linux systems:  
`export MQSERVER='ChannelName/TransportType/ConnectionName'`
-  For IBM i:  
`ADDENVVAR ENVVAR(MQSERVER) VALUE('ChannelName/TransportType/ConnectionName')`
-  For z/OS  
`export MQSERVER='SYSTEM.DEF.SVRCONN/TCP/AMACHINE.ACOMPANY.COM(1414) '`

*TransportType* can be one of the following values, depending on your IBM MQ client platform:

- LU62
- TCP
- NETBIOS
- SPX

*ConnectionName* can be a comma-separated list of connection names. The connection names in the list are used in a similar way to multiple connections in a client connection table. The *ConnectionName* list might be used as an alternative to queue manager groups to specify multiple connections for the client to try. If you are configuring a multi-instance queue manager, you might use a *ConnectionName* list to specify different queue manager instances.

#### **TCP/IP default port:**

By default, for TCP/IP, IBM MQ assumes that the channel will be connected to port 1414.

You can change this by:

- Adding the port number in brackets as the last part of the *ConnectionName*:
  - For Windows:
 

```
SET MQSERVER=ChannelName/TransportType/ConnectionName(PortNumber)
```
  - For UNIX and Linux systems:
 

```
export MQSERVER='ChannelName/TransportType/ConnectionName(PortNumber)'
```
- Changing the `mqclient.ini` file by adding the port number to the protocol name, for example:
 

```
TCP:
port=2001
```
- Adding IBM MQ to the services file as described in “Using the TCP/IP listener” on page 874.

#### **SPX default socket:**

By default, for SPX, IBM MQ assumes that the channel will be connected to socket 5E86.

You can change this by:

- Adding the socket number in brackets as the last part of the *ConnectionName*:
 

```
SET MQSERVER=ChannelName/TransportType/ConnectionName(SocketNumber)
```

For SPX connections, specify the *ConnectionName* and socket in the form `network.node(socket)`. If the IBM MQ client and server are on the same network, the network need not be specified. If you are using the default socket, the socket need not be specified.
- Changing the `qm.ini` file by adding the port number to the protocol name, for example:
 

```
SPX:
socket=5E87
```



## Using MQSERVER:

If you use the MQSERVER environment variable to define the channel between your IBM MQ MQI client machine and a server machine, this is the only channel available to your application, and no reference is made to the client channel definition table (CCDT).

In this situation, the listener program that you have running on the server machine determines the queue manager to which your application will connect. It will be the same queue manager as the listener program is connected to.

If the MQCONN or MQCONNX request specifies a queue manager other than the one the listener is connected to, or if the MQSERVER parameter *TransportType* is not recognized, the MQCONN or MQCONNX request fails with return code MQRC\_Q\_MGR\_NAME\_ERROR.

On UNIX and Linux systems, you might define MQSERVER as in one of the following examples:

```
export MQSERVER=CHANNEL1/TCP/'9.20.4.56(2002)'
export MQSERVER=CHANNEL1/LU62/BOX99
```

All MQCONN or MQCONNX requests then attempt to use the channel you have defined unless an MQCD structure has been referenced from the MQCNO structure supplied to MQCONNX, in which case the channel specified by the MQCD structure takes priority over any specified by the MQSERVER environment variable.

The MQSERVER environment variable takes priority over any client channel definition pointed to by MQCHLLIB and MQCHLTAB.

## Canceling MQSERVER

To cancel MQSERVER and return to the client channel definition table pointed to by MQCHLLIB and MQCHLTAB, enter the following:

- On Windows:  
SET MQSERVER=
- On UNIX and Linux systems:  
unset MQSERVER

## MQSSLCRYP

MQSSLCRYP holds a parameter string that allows you to configure the cryptographic hardware present on the system. The permitted values are the same as for the SSLCRYP parameter of the ALTER QMGR command.

To set this variable use one of these commands:


- On Windows systems:  
SET MQSSLCRYP=string
- On UNIX and Linux systems:  
export MQSSLCRYP=string

## MQSSLFIPS

MQSSLFIPS specifies whether only FIPS-certified algorithms are to be used if cryptography is carried out in IBM MQ. The values are the same as for the SSLFIPS parameter of the ALTER QMGR command.

The use of FIPS-certified algorithms is affected by the use of cryptographic hardware, see Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client.

To set this variable use one of these commands:

- On Windows systems:  
SET MQSSLFIPS=YES|NO
- On UNIX and Linux systems:  
export MQSSLFIPS=YES|NO
-  On IBM i:  
ADDENVVAR ENVVAR(MQSSLFIPS) VALUE(YES|NO)


The default is NO.

## MQSSLKEYR

MQSSLKEYR specifies the location of the key repository that holds the digital certificate belonging to the user, in stem format. Stem format means that it includes the full path and the file name without an extension.

For full details, see the SSLKEYR parameter of the ALTER QMGR command.

To set this variable use one of these commands:

- On Windows systems:  
SET MQSSLKEYR=pathname
- On UNIX and Linux systems:  
export MQSSLKEYR=pathname
-  On IBM i:  
ADDENVVAR ENVVAR(MQSSLKEYR) VALUE(pathname)

There is no default value.

## MQSSLPROXY

MQSSLPROXY specifies the host name and port number of the HTTP proxy server to be used by GSKit for OCSF checks.

To set this variable use one of these commands:

- On Windows systems:  
SET MQSSLPROXY= *string*
- On UNIX and Linux systems:  
export MQSSLPROXY="*string*"

The string is either the host name or network address of the HTTP Proxy server which is to be used by GSKit for OCSF checks. This address can be followed by an optional port number, enclosed in parentheses. If you do not specify the port number, the default HTTP port, 80, is used.

For example, on UNIX and Linux systems, you can use the one of the following commands:

- export MQSSLPROXY="proxy.example.com(80)"
- export MQSSLPROXY="127.0.0.1"


## MQSSLRESET

MQSSLRESET represents the number of unencrypted bytes sent and received on an SSL or TLS channel before the secret key is renegotiated.

See Resetting SSL and TLS secret keys for more information about secret key renegotiation.

It can be set to an integer in the range 0 through 999 999 999. The default is 0, which indicates that secret keys are never renegotiated. If you specify an SSL or TLS secret key reset count in the range 1 byte through 32 KB, SSL or TLS channels use a secret key reset count of 32 KB. This secret reset count is to avoid excessive key resets which would occur for small SSL or TLS secret key reset values.

To set this variable use one of these commands:

- On Windows systems:  
SET MQSSLRESET=integer
- On UNIX and Linux systems:  
export MQSSLRESET=integer
-  On IBM i:  
ADDENVVAR ENVVAR(MQSSLRESET) VALUE(integer)

## MQSUITEB

You can configure IBM MQ to operate in compliance with the NSA Suite B standard on Windows, UNIX and Linux platforms.

Suite B restricts the set of enabled cryptographic algorithms in order to provide an assured level of security.

See Configuring IBM MQ for Suite B for more information.

## MQTCPTIMEOUT

How long IBM MQ waits for a TCP connect call.

---

## Changing IBM MQ and queue manager configuration information

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

You can change IBM MQ configuration information by changing the values specified on a set of configuration attributes (or parameters) that govern IBM MQ.

Change attribute information by editing the IBM MQ configuration files. On IBM MQ for Windows and Linux (x86 and x86-64 platforms), the IBM MQ configuration files can be edited using the MQ Explorer.


On Windows systems you can also use amqmdain to change configuration information, as described in amqmdain

To find out more about configuring IBM MQ and queue managers for your platform, see the following subtopics:

### Related concepts:

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

 “Configuring queue managers on z/OS” on page 1194

Use these instructions to configure queue managers on IBM MQ for z/OS.

### Related information:

Planning

Administering IBM MQ

## Changing configuration information on Windows, UNIX and Linux systems

Configuration attributes are held in configuration files, at the level of the node and of the queue manager.

On Windows, UNIX and Linux platforms, you can change IBM MQ configuration attributes within:

- An IBM MQ configuration file ( **mqs.ini** ) to effect changes for IBM MQ on the node as a whole. There is one mqs.ini file for each node.

See “Attributes for changing IBM MQ configuration information” on page 775 for more information on the stanzas included in **mqs.ini**.

- A queue manager configuration file ( **qm.ini** ) to effect changes for specific queue managers. There is one qm.ini file for each queue manager on the node.

See “Changing queue manager configuration information” on page 782 for more information on the stanzas included in **qm.ini**.

Client configuration options are held separately, in the client configuration file.

A configuration file (or *stanza* file) contains one or more stanzas, which are groups of lines in the .ini file that together have a common function or define part of a system, such as log functions, channel functions, and installable services.

Because the IBM MQ configuration file is used to locate the data associated with queue managers, a nonexistent or incorrect configuration file can cause some or all MQSC commands to fail. Also, applications cannot connect to a queue manager that is not defined in the IBM MQ configuration file.

Any changes you make to a configuration file usually do not take effect until the next time the queue manager is started.

On Windows and Linux (x86 and x86-64 platforms) systems, you can edit configuration information from the MQ Explorer.

On Windows systems you can also use the the **amqmdain** command to edit the configuration files.

For more information about the configuration options on Windows, UNIX and Linux systems, see the following subtopics:

**Related concepts:**


“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

**Related reference:**

 “Changing configuration information on IBM i” on page 763

Use this information to learn how to change the behavior of queue managers to suit your installation's needs.

“Attributes for changing IBM MQ configuration information” on page 775

On IBM MQ for Windows systems and on IBM MQ for Linux (x86 and x86-64 platforms) systems, modify configuration information using the IBM MQ Explorer. On other systems, modify the information by editing the mq.ini configuration file.

“Changing queue manager configuration information” on page 782

The attributes described here modify the configuration of an individual queue manager. They override any settings for IBM MQ.

**Related information:**

Planning

Administering IBM MQ

**Editing configuration files**

Edit configuration files using commands or a standard text editor.

Before editing a configuration file, back it up so that you have a copy you can revert to if the need arises.

You can edit configuration files either:

- Automatically, using commands that change the configuration of queue managers on the node
- Manually, using a standard text editor

You can edit the default values in the IBM MQ configuration files after installation.

If you set an incorrect value on a configuration file attribute, the value is ignored and an operator message is issued to indicate the problem. (The effect is the same as missing out the attribute entirely.)

When you create a new queue manager:

- Back up the IBM MQ configuration file
- Back up the new queue manager configuration file

Comments can be included in configuration files by adding a “;” or a “#” character before the comment text. If you want to use a “;” or a “#” character without it representing a comment, you can prefix the character with a “\” character and it will be used as part of the configuration data.

**When do you need to edit a configuration file?**

Edit a configuration file to recover from backup, move a queue manager, change the default queue manager or to assist IBM support.

You might need to edit a configuration file if, for example:

- You lose a configuration file. (Recover from backup if you can.)
- You need to move one or more queue managers to a new directory.

- You need to change your default queue manager; this could happen if you accidentally delete the existing queue manager.
- You are advised to do so by your IBM Support Center.

## Configuration file priorities

The value of an attribute is defined in multiple places. Attributes set in commands take precedence over attributes in configuration files.

The attribute values of a configuration file are set according to the following priorities:

- Parameters entered on the command line take precedence over values defined in the configuration files
- Values defined in the `qm.ini` files take precedence over values defined in the `mqs.ini` file

## The IBM MQ configuration file, `mqs.ini`

The IBM MQ configuration file, `mqs.ini`, contains information relevant to all the queue managers on the node. It is created automatically during installation.

On IBM MQ for UNIX and Linux products, the data directory and log directory are always `/var/mqm` and `/var/mqm/log` respectively.

On Windows systems, the location of the data directory `mqs.ini`, and the location of the log directory, are stored in the registry, as their location can vary.

In addition, on Windows systems, the installation configuration information (contained in `mqinst.ini` on IBM MQ for UNIX and Linux systems) is in the registry, as there is no `mqinst.ini` file on Windows.

The `mqs.ini` file for Windows systems is given by the `WorkPath` specified in the `HKLM\SOFTWARE\IBM\WebSphere MQ` key. It contains:

- The names of the queue managers
- The name of the default queue manager
- The location of the files associated with each of them

The supplied `LogDefaults` stanza for a new IBM MQ installation does not contain any explicit values for the attributes. The lack of an attribute means that the default for this value is used upon creation of a new queue manager. The default values are shown for the `LogDefaults` stanza in Figure 82 on page 759. A value of zero for the `LogBufferPages` attribute means 512.

If you require a non-default value, you must explicitly specify that value in the `LogDefaults` stanza.

```

*****#
** Module Name: mqs.ini **
** Type : IBM MQ Machine-wide Configuration File **
** Function : Define IBM MQ resources for an entire machine **
*****#
** Notes : **
** 1) This is the installation time default configuration **
** **
*****#
AllQueueManagers:
*****#
** The path to the qmgrs directory, below which queue manager data **
** is stored **
*****#
DefaultPrefix=/var/mqm

LogDefaults:
 LogPrimaryFiles=3
 LogSecondaryFiles=2
 LogFilePages=4096
 LogType=CIRCULAR
 LogBufferPages=0
 LogDefaultPath=/var/mqm/log

QueueManager:
 Name=saturn.queue.manager
 Prefix=/var/mqm
 Directory=saturn!queue!manager
 InstallationName=Installation1

QueueManager:
 Name=pluto.queue.manager
 Prefix=/var/mqm
 Directory=pluto!queue!manager
 InstallationName=Installation2

DefaultQueueManager:
 Name=saturn.queue.manager

ApiExitTemplate:
 Name=OurPayrollQueueAuditor
 Sequence=2
 Function=EntryPoint
 Module=/usr/ABC/auditor
 Data=123

ApiExitCommon:
 Name=MQPoliceman
 Sequence=1
 Function=EntryPoint
 Module=/usr/MQPolice/tmqp
 Data=CheckEverything

```

Figure 82. Example of an IBM MQ configuration file for UNIX systems

## Queue manager configuration files, qm.ini

A queue manager configuration file, qm.ini, contains information relevant to a specific queue manager.

There is one queue manager configuration file for each queue manager. The qm.ini file is automatically created when the queue manager with which it is associated is created.

**V 8.0.0.8** From IBM MQ Version 8.0.0, Fix Pack 8, the **strmqm** command checks the syntax of the CHANNELS and SSL stanzas in the `qm.ini` file before starting the queue manager fully, which makes it much easier to see what is wrong, and correct it quickly if **strmqm** finds that the `qm.ini` file contains any errors. For more information, see `strmqm`.

On UNIX and Linux systems a `qm.ini` file is held in the root of the directory tree occupied by the queue manager. For example, the path and the name for a configuration file for a queue manager called QMNAME is:

```
/var/mqm/qmgrs/QMNAME/qm.ini
```

On Windows systems the location of the `qm.ini` file is given by the `WorkPath` specified in the `HKLM\SOFTWARE\IBM\WebSphere MQ` key. For example, the path and the name for a configuration file for a queue manager called QMNAME is:

```
C:\ProgramData\IBM\MQ\qmgrs\QMNAME\qm.ini
```

The queue manager name can be up to 48 characters in length. However, this does not guarantee that the name is valid or unique. Therefore, a directory name is generated based on the queue manager name. This process is known as *name transformation*. For a description, see *Understanding IBM MQ file names*.

Figure 83 on page 761 shows how groups of attributes might be arranged in a queue manager configuration file in IBM MQ for UNIX and Linux systems.



```

** Module Name: qm.ini **
** Type : IBM MQ queue manager configuration file **
Function : Define the configuration of a single queue manager **
** **
*****#
** Notes : **
** 1) This file defines the configuration of the queue manager **
** **
*****#

ExitPath:
 ExitsDefaultPath=/var/mqm/exits
 ExitsDefaultPath64=/var/mqm/exits64

Service:
 Name=AuthorizationService
 EntryPoints=14

ServiceComponent:
 Service=AuthorizationService
 Name=MQSeries.UNIX.auth.service
 Module=amqzfu
 ComponentDataSize=0

Log:
 LogPrimaryFiles=3
 LogSecondaryFiles=2
 LogFilePages=4096
 LogType=CIRCULAR
 LogBufferPages=0 1
 LogPath=/var/mqm/log/saturn!queue!manager/

XAResourceManager:
 Name=DB2 Resource Manager Bank
 SwitchFile=/usr/bin/db2swit
 XAOpenString=MQBankDB
 XACloseString=
 ThreadOfControl=THREAD

Channels: 2
 MaxChannels=200
 MaxActiveChannels=100
 MQIBindType=STANDARD

TCP:
 SndBuffSize=0
 RcvBuffSize=0
 RcvSndBuffSize=0
 RcvRcvBuffSize=0
 ClntSndBuffSize=0
 ClntRcvBuffSize=0
 SvrSndBuffSize=0
 SvrRcvBuffSize=0

QMErrorLog:
 ErrorLogSize=262144
 ExcludeMessage=7234
 SuppressMessage=9001,9002,9202
 SuppressInterval=30

ApiExitLocal:
 Name=ClientApplicationAPIchecker
 Sequence=3
 Function=EntryPoint
 Module=/usr/Dev/ClientAppChecker
 Data=9.20.176.20

```

Figure 83. Example queue manager configuration file for IBM MQ for UNIX and Linux systems

Note that the `qm.ini` file for Windows includes an additional *AccessMode* stanza:

```
AccessMode:
SecurityGroup=wmq\wmq
```

Notes for Figure 83 on page 761:

1. The value of zero for `LogBufferPages` gives a value of 512.
2. For more information on the Channel stanza, see “Initialization and configuration files” on page 852.
3. The maximum number of `XAResourceManager` stanzas is limited to 255. However, you should use only a small number of stanzas to avoid transaction performance degradation.

See “Changing configuration information on Windows, UNIX and Linux systems” on page 756 for information on when your changes take effect.

## Installation configuration file, `mqinst.ini` UNIX and Linux systems

The installation configuration file, `mqinst.ini`, contains information about all the IBM MQ installations on a UNIX or Linux system.

The `mqinst.ini` file is in the `/etc/opt/mqm` directory on UNIX and Linux systems. It contains information about which installation, if any, is the primary installation as well as the following information for each installation:

- The installation name
- The installation description
- The installation identifier
- The installation path

This file must not be edited or referenced directly since its format is not fixed, and could change. Instead, use the following commands to create, delete, query, and modify, the values in the `mqinst.ini` file:

`crtmqinst` to create entries.

`dltmqinst` to delete entries.

`dspmqinst` to display entries.

`setmqinst` to set entries.

The installation identifier, for internal use only, is set automatically and must not be changed.

## Windows systems

Installation configuration information is held in the following key on Windows systems:

```
HKLM\SOFTWARE\IBM\WebSphere MQ\Installation\<InstallationName>
```

This key must not be edited or referenced directly since its format is not fixed, and could change. Instead, use the following commands to query, and modify, the values in the registry:

`dspmqinst` to display entries.

`setmqinst` to set entries.

On Windows, the **`crtmqinst`** and **`dltmqinst`** commands are not available. The installation and uninstallation processes handle the creation and deletion of the required registry entries.

## Changing configuration information on IBM i

Use this information to learn how to change the behavior of queue managers to suit your installation's needs.

You change IBM MQ configuration information by modifying the values specified on a set of configuration attributes (or parameters) that govern IBM MQ. You change these attributes by editing the **IBM MQ configuration files**.

For information on modifying the configuration values on IBM i, see the following topics:

- “IBM MQ for IBM i configuration files”
- “Attributes for changing IBM MQ for IBM i configuration information” on page 765
- “Changing IBM MQ queue manager configuration information” on page 767
- “Example IBM i mqs.ini and qm.ini files” on page 772

### Related concepts:

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

“Changing configuration information on Windows, UNIX and Linux systems” on page 756

Configuration attributes are held in configuration files, at the level of the node and of the queue manager.

### Related reference:

“Attributes for changing IBM MQ configuration information” on page 775

On IBM MQ for Windows systems and on IBM MQ for Linux (x86 and x86-64 platforms) systems, modify configuration information using the IBM MQ Explorer. On other systems, modify the information by editing the mqs.ini configuration file.

“Changing queue manager configuration information” on page 782

The attributes described here modify the configuration of an individual queue manager. They override any settings for IBM MQ.

### Related information:

Planning

Administering IBM MQ

## IBM MQ for IBM i configuration files

Use this information to understand the methods for configuring IBM MQ for IBM i.

You modify IBM MQ configuration attributes within:

- An IBM MQ configuration file, `mqs.ini`, effect changes on the node as a whole. There is one `mqs.ini` file for each IBM MQ installation.
- A queue manager configuration file, `qm.ini`, effect changes for specific queue managers. There is one `qm.ini` file for each queue manager on the node.

Note that `.ini` files are stream files resident in the IFS.

A configuration file (which can be referred to as a *stanza* file) contains one or more stanzas, which are groups of lines in the `.ini` file that together have a common function or define part of a system, for example, log functions and channel functions. Any changes made to a configuration file do not take effect until the next time the queue manager is started.

## Editing configuration files

Before editing a configuration file, back it up so that you have a copy you can revert to if the need arises.

You can edit configuration files either:

- Automatically, using commands that change the configuration of queue managers on the node.
- Manually, using the EDTF CL editor.

You can edit the default values in the IBM MQ configuration files after installation. If you set an incorrect value on a configuration file attribute, the value is ignored and an operator message is issued to indicate the problem. (The effect is the same as missing out the attribute entirely.)

When you create a new queue manager:

- Back up the IBM MQ configuration file.
- Back up the new queue manager configuration file.

## When do you need to edit a configuration file?

You might need to edit a configuration file if, for example:

- You lose a configuration file; recover from backup if possible.
- You need to move one or more queue managers to a new directory.
- You need to change your default queue manager; this could happen if you accidentally delete the existing queue manager.
- You are advised to do so by your IBM Support Center.

## Configuration file priorities

The attribute values of a configuration file are set according to the following priorities:

- Parameters entered on the command line take precedence over values defined in the configuration files.
- Values defined in the `qm.ini` files take precedence over values defined in the `mqs.ini` file.

## The IBM MQ configuration file `mqs.ini`

The IBM MQ configuration file, `mqs.ini`, contains information relevant to all the queue managers on an IBM MQ installation. It is created automatically during installation. In particular, the `mqs.ini` file is used to locate the data associated with each queue manager.

The `mqs.ini` file is stored in `/QIBM/UserData/mqm`

The `mqs.ini` file contains:

- The names of the queue managers.
- The name of the default queue manager.
- The location of the files associated with each queue manager.
- Information identifying any API exits (see [Configuring API exits](#) for more information).

## Queue manager configuration files `qm.ini`

A queue manager configuration file, `qm.ini`, contains information relevant to a specific queue manager. There is one queue manager configuration file for each queue manager. The `qm.ini` file is automatically created when the queue manager with which it is associated is created.

A `qm.ini` file is held in the `<mqmdata directory>/QMNAME/qm.ini`, where `<mqmdata directory>` is `/QIBM/UserData/mqm` by default and `QMNAME` is the name of the queue manager to which the initialization file applies.

**Note:**

1. You can change the `<mqmdata directory>` in the `mqs.ini` file.
2. The queue manager name can be up to 48 characters in length. However, this does not guarantee that the name is valid or unique. Therefore, a directory name is generated based on the queue manager name. This process is known as **name transformation**. See *Understanding IBM MQ for IBM i queue manager library names* for further information.

## Attributes for changing IBM MQ for IBM i configuration information

Use this information to understand the configuration information stanzas.

The following groups of attributes occur in `mqs.ini`:

- “The `AllQueueManagers` stanza”
- “The `DefaultQueueManager` stanza” on page 766
- “The `ExitProperties` stanza” on page 766
- “The `QueueManager` stanza” on page 767

There are also two stanzas associated with API exits, `ApiExitCommon` and `ApiExitTemplate`. For details on using these, see *Configuring API exits*.

### The `AllQueueManagers` stanza

The `AllQueueManagers` stanza can specify:

- The path to the `qmgrs` directory where the files associated with a queue manager are stored
- The path to the executable library
- The method for converting EBCDIC-format data to ASCII format

In the descriptions of the stanzas, the value underlined is the default value and the `|` symbol means *or*.

**DefaultPrefix=** *directory\_name*

The path to the `qmgrs` directory, within which the queue manager data is kept. If you change the default prefix for the queue manager, you must replicate the directory structure that was created at installation time. In particular, you must create the `qmgrs` structure. Stop IBM MQ before changing the default prefix, and restart IBM MQ only after moving the structures to the new location and changing the default prefix.

As an alternative to changing the default prefix, you can use the environment variable `MQSPREFIX` to override the `DefaultPrefix` for the **CRTMQM** command.

**ConvEBCDICNewline=**`NL_TO_LF`**|**`TABLE`**|**`ISO`

EBCDIC code pages contain a newline (NL) character that is not supported by ASCII code pages, although some ISO variants of ASCII contain an equivalent.

Use the `ConvEBCDICNewline` attribute to specify the method IBM MQ is to use when converting the EBCDIC NL character into ASCII format.

**NL\_TO\_LF**

Convert the EBCDIC NL character (X'15') to the ASCII line feed character, LF (X'0A'), for all EBCDIC to ASCII conversions.

`NL_TO_LF` is the default.

**TABLE**

Convert the EBCDIC NL character according to the conversion tables used on IBM i for all EBCDIC to ASCII conversions.

Note that the effect of this type of conversion can vary from language to language .

## ISO

Specify ISO if you want:

- ISO CCSIDs to be converted using the TABLE method
- All other CCSIDs to be converted using the NL\_TO\_CF method.

Possible ISO CCSIDs are shown in Table 120.

Table 120. List of possible ISO CCSIDs

| CCSID | Code Set  |
|-------|-----------|
| 819   | ISO8859-1 |
| 912   | ISO8859-2 |
| 915   | ISO8859-5 |
| 1089  | ISO8859-6 |
| 813   | ISO8859-7 |
| 916   | ISO8859-8 |
| 920   | ISO8859-9 |
| 1051  | roman8    |

If the ASCII CCSID is not an ISO subset, ConvEBCDICNewline defaults to NL\_TO\_LF.

## The DefaultQueueManager stanza

The DefaultQueueManager stanza specifies the default queue manager for the node.

**Name=** *default\_queue\_manager*

The default queue manager processes any commands for which a queue manager name is not explicitly specified. The DefaultQueueManager attribute is automatically updated if you create a new default queue manager. If you inadvertently create a new default queue manager and then want to revert to the original, you must alter the DefaultQueueManager attribute manually.

## The ExitProperties stanza

The ExitProperties stanza specifies configuration options used by queue manager exit programs.

In the descriptions of the stanzas, the value underlined is the default value and the | symbol means *or*.

**CLWLMode=** SAFE | FAST

The cluster workload exit, CLWL, allows you to specify which cluster queue in the cluster is to be opened in response to an MQI call (for example: MQOPEN or MQPUT). The CLWL exit runs either in FAST mode or SAFE mode depending on the value you specify on the CLWLMode attribute. If you omit the CLWLMode attribute, the cluster workload exit runs in SAFE mode.

### SAFE

Run the CLWL exit in a separate process to the queue manager. This is the default.

If a problem arises with the user-written CLWL exit when running in SAFE mode, the following happens:

- The CLWL server process (amqzlw0) fails
- The queue manager restarts the CLWL server process
- The error is reported to you in the error log. If an MQI call is in progress, you receive notification in the form of a bad return code.

The integrity of the queue manager is preserved.

**Note:** Running the CLWL exit in a separate process might have a detrimental effect on performance.

#### **FAST**

Run the cluster exit inline in the queue manager process.

Specifying this option improves performance by avoiding the overheads associated with running in SAFE mode, but does so at the expense of queue manager integrity. Run the CLWL exit in FAST mode only if you are convinced that there are **no** problems with your CLWL exit, and you are particularly concerned about performance overheads.

If a problem arises when the CLWL exit is running in FAST mode, the queue manager fails and you run the risk of compromising the integrity of the queue manager.

### **The QueueManager stanza**

There is one QueueManager stanza for every queue manager. These attributes specify the queue manager name and the name of the directory containing the files associated with that queue manager. The name of the directory is based on the queue manager name, but is transformed if the queue manager name is not a valid file name.

See Understanding IBM MQ for IBM i queue manager library names for more information about name transformation.

**Name=** *queue\_manager\_name*

The name of the queue manager.

**Prefix=** *prefix*

Where the queue manager files are stored. By default, this is the same as the value specified on the DefaultPrefix attribute of the AllQueueManager stanza in the mq5.ini file.

**Directory=** *name*

The name of the subdirectory under the <prefix>\QMGRS directory where the queue manager files are stored. This name is based on the queue manager name, but can be transformed if there is a duplicate name, or if the queue manager name is not a valid file name.

**Library=** *name*

The name of the library where IBM i objects pertinent to this queue manager, for example, journals and journal receivers, are stored. This name is based on the queue manager name, but can be transformed if there is a duplicate name, or if the queue manager name is not a valid library name.

### **Changing IBM MQ queue manager configuration information**

Use this information to understand the queue manager configuration stanzas.

There are two stanzas associated with API exits, ApiExitCommon and ApiExitTemplate. For details on using these, see Configuring API exits.

The following groups of attributes can occur in a qm.ini file for a specific queue manager, or used to override values set in mq5.ini.

See the following topics for changing configuration information for specific options:

- “The Log stanza” on page 768
- “The Channels stanza” on page 768
- “The queue manager error log stanza” on page 770
- “The TCP stanza” on page 771
- PreConnect stanza of the client configuration file

## The Log stanza:

Parameters for configuring the log file.

The Log stanza specifies the log attributes for a particular queue manager. By default, these are inherited from the settings specified in the LogDefaults stanza in the mq.ini file when the queue manager is created.

Only change attributes of this stanza if you want to configure a queue manager differently from others.

The values specified on the attributes in the qm.ini file are read when the queue manager is started. The file is created when the queue manager is created.

### LogBufferSize

The journal buffer size, in bytes. Enter a number in the range 32 000 through 15 761 440. The default is 32 000.

### LogPath= *library\_name*

The name of the library used to store journals and journal receivers for this queue manager.

### LogReceiverSize

The journal receiver size, in kilobytes. The default is 100 000.

## The Channels stanza:

The Channels stanza contains information about the channels.

### MaxChannels= 100 | *number*

The maximum number of *current* channels allowed. For z/OS, the value must be 1 - 9999, with a default value of 200. For all other platforms, the value must be 1 - 65,535, with a default value of 100.

### MaxActiveChannels= *MaxChannels\_value*

The maximum number of channels allowed to be *active* at any time. The default is the value specified on the MaxChannels attribute.

### MaxInitiators= 3 | *number*

The maximum number of initiators. The default and maximum value is 3.

### MQIBINDTYPE=FASTPATH | STANDARD

The binding for applications.

#### FASTPATH

Channels connect using MQCONNX FASTPATH. That is, there is no agent process.

#### STANDARD

Channels connect using STANDARD.

### ThreadedListener= NO | YES

Whether to start RUNMQLSR ( YES ) or AMQCLMAA ( NO ) as a listener.

If you specify ThreadedListener=YES, all channels run as threads of a single job. This limits the number of connections to the resources available to a single job.

If you specify ThreadedListener=NO, the non-threaded listener (AMQCLMAA) starts a new responder job (AMQCRSTA) for each inbound TCP/IP channel. The disadvantage of this technique is that it is not as fast to start a new AMQCRSTA job as it is to start a thread within a RUNMQLSR job, therefore connection times for a non-threaded listener are slower than those for a threaded listener.

### AdoptNewMCA= NO | SVR | SNDR | RCVR | CLUSRCVR | ALL | FASTPATH

If IBM MQ receives a request to start a channel, but finds that an amqcrsta process exists for the same channel, the existing process must be stopped before the new one can start. The AdoptNewMCA attribute allows you to control the ending of an existing process and the startup of a new one for a specified channel type.



If you specify the `AdoptNewMCA` attribute for a given channel type, but the new channel fails to start because the channel is already running:

1. The new channel tries to end the previous one.
2. If the previous channel server does not end by the time the `AdoptNewMCATimeout` wait interval expires, the process (or the thread) for the previous channel server is ended.
3. If the previous channel server has not ended after step 2, and after the `AdoptNewMCATimeout` wait interval expires for a second time, IBM MQ ends the channel with a `CHANNEL IN USE` error.

You specify one or more values, separated by commas or blanks, from the following list:

**NO** The `AdoptNewMCA` feature is not required. This is the default.

**SVR**

Adopt server channels

**SNDP**

Adopt sender channels

**RCVR**

Adopt receiver channels

**CLUSRCVR**

Adopt cluster receiver channels

**ALL**

Adopt all channel types, except for `FASTPATH` channels

**FASTPATH**

Adopt the channel if it is a `FASTPATH` channel. This happens only if the appropriate channel type is also specified, for example, `AdoptNewMCA=RCVR,SVR,FASTPATH`

**Attention!:** The `AdoptNewMCA` attribute can behave in an unpredictable fashion with `FASTPATH` channels because of the internal design of the queue manager. Exercise great caution when enabling the `AdoptNewMCA` attribute for `FASTPATH` channels.

**AdoptNewMCATimeout= 60 | 1-3600**

The amount of time, in seconds, that the new process waits for the old process to end. Specify a value, in seconds, in the range 1 - 3600. The default value is 60.

**AdoptNewMCACheck=QM|ADDRESS|NAME|ALL**

The `AdoptNewMCACheck` attribute allows you to specify the type checking required when enabling the `AdoptNewMCA` attribute. It is important for you to perform all three of the following checks, if possible, to protect your channels from being shut down, inadvertently or maliciously. At the very least check that the channel names match.

Specify one or more values, separated by commas or blanks, from the following:

**QM** The listener process checks that the queue manager names match.

**ADDRESS**

The listener process checks the communications address, for example, the TCP/IP address.

**NAME**

The listener process checks that the channel names match.

**ALL**

The listener process checks for matching queue manager names, the communications address, and for matching channel names.

The default is `AdoptNewMCACheck=NAME,ADDRESS,QM`.

**Related concepts:**

“Channel states” on page 834

A channel can be in one of many states at any time. Some states also have substates. From a given state a channel can move into other states.

**The queue manager error log stanza:**

Use the QMErrorLog stanza in the `qm.ini` file to tailor the operation and contents of queue manager error logs.

**ErrorLogSize=** *maxsize*

Specifies the size of the queue manager error log at which it is copied to the backup. *maxsize* must be in the range 1048576 through 2147483648 bytes. If **ErrorLogSize** is not specified, the default value of 2097152 bytes (2 MB) is used.

**ExcludeMessage=** *msgIds*

Specifies messages that are not to be written to the queue manager error log. *msgIds* contain a comma separated list of message IDs from the following:

- 7163 - Job started message ( IBM i only)
- 7234 - Number of messages loaded
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel
- 9558 - Remote Channel is not available
- 9776 - Channel was blocked by user id
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address
- 9999 - Channel program ended abnormally

**SuppressMessage=** *msgIds*

Specifies messages that will be written to the queue manager error log once only in a specified time interval. The time interval is specified by **SuppressInterval** . *msgIds* contain a comma separated list of message IDs from the following:

- 7163 - Job started message ( IBM i only)
- 7234 - Number of messages loaded
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel

- 9558 - Remote Channel is not available
- 9776 - Channel was blocked by user id
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address
- 9999 - Channel program ended abnormally

If the same message id is specified in both **SuppressMessage** and **ExcludeMessage**, the message is excluded.

**SuppressInterval=** *length*

Specifies the time interval, in seconds, in which messages specified in **SuppressMessage** will be written to the queue manager error log once only. *length* must be in the range 1 through 86400 seconds. If **SuppressInterval** is not specified, the default value of 30 seconds is used.

**The TCP stanza:**

Use these queue manager properties pages, or stanzas in the `qm.ini` file, to specify network protocol configuration parameters. They override the default attributes for channels.

**Note:** Only attributes representing changes to the default values need to be specified.

**TCP**

The following attributes can be specified:

**Port=** 1414 | *port\_number*

The default port number, in decimal notation, for TCP/IP sessions. The default port number for IBM MQ Version 8.0 is 1414.

**KeepAlive=** NO | YES

Switch the KeepAlive function on or off. KeepAlive=YES causes TCP/IP to check periodically that the other end of the connection is still available. If it is not, the channel is closed.

**ListenerBacklog=***number*

When receiving on TCP/IP, a maximum number of outstanding connection requests is set. This can be considered to be a *backlog* of requests waiting on the TCP/IP port for the listener to accept the request. The default listener backlog value for IBM i is 255; the maximum is 512. If the backlog reaches the value of 512, the TCP/IP connection is rejected and the channel cannot start.

For MCA channels, this results in the channel going into a RETRY state and retrying the connection at a later time.

For client connections, the client receives an MQRC\_Q\_MGR\_NOT\_AVAILABLE reason code from MQCONN and should retry the connection at a later time.

The ListenerBacklog attribute allows you to override the default number of outstanding requests for the TCP/IP listener.

**Connect\_Timeout=***number* | 0

The number of seconds before an attempt to connect the socket times out. The default value of zero specifies that there is no connect timeout.

The following group of properties can be used to control the size of buffers used by TCP/IP. The values are passed directly to the TCP/IP layer of the operating system. Great care should be taken when using these properties. If the values are set incorrectly it can adversely affect the TCP/IP performance. For further information about how this affects performance refer to the TCP/IP documentation for your environment. A value of zero indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by IBM MQ.

**SndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the sending end of channels. This stanza value can be overridden by a stanza more specific to the channel type, for example RcvSndBuffSize.

**RcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the receiving end of channels. This stanza value can be overridden by a stanza more specific to the channel type, for example RcvRcvBuffSize. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**RcvSndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the sender end of a receiver channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**RcvRcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the receiving end of a receiver channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**SvrSndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the server end of a client-connection server-connection channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**SvrRcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the server end of a client-connection server-connection channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

## Example IBM i mqs.ini and qm.ini files

Figure 84 on page 773 shows an example of an mqs.ini file for IBM i.

```

#####
#* Module Name: mqs.ini *#
#* Type : IBM MQ Configuration File *#
#* Function : Define IBM MQ resources for the node *#
#* *#
#####
#* Notes : *#
#* 1) This is an example IBM MQ configuration file *#
#* *#
#####
AllQueueManagers:
#####
#* The path to the qmgrs directory, within which queue manager data *#
#* is stored *#
#####
DefaultPrefix=/QIBM/UserData/mqm

QueueManager:
Name=saturn.queue.manager
Prefix=/QIBM/UserData/mqm
Library=QMSATURN.Q
Directory=saturn!queue!manager

QueueManager:
Name=pluto.queue.manager
Prefix=/QIBM/UserData/mqm
Library=QMPLUTO.QU
Directory=pluto!queue!manager

DefaultQueueManager:
Name=saturn.queue.manager

```

*Figure 84. Example of an IBM MQ configuration file*

Figure 85 on page 774 shows how groups of attributes might be arranged in a queue manager configuration file for IBM i.

```

#####
#* Module Name: qm.ini *#
#* Type : IBM MQ queue manager configuration file *#
Function : Define the configuration of a single queue manager *#
#* *#
#####
#* Notes : *#
#* 1) This file defines the configuration of the queue manager *#
#* *#
#####
Log:
LogPath=QMSATURN.Q
LogReceiverSize=65536

CHANNELS:
MaxChannels = 20 ; Maximum number of channels allowed.
 ; Default is 100.
MaxActiveChannels = 10 ; Maximum number of channels allowed to be
 ; active at any time. The default is the
 ; value of MaxChannels.

TCP:
KeepAlive = Yes ; TCP/IP entries.
 ; Switch KeepAlive on.
SvrSndBuffSize=20000 ; Size in bytes of the TCP/IP send buffer for each
 ; channel instance. Default is 32768.
SvrRcvBuffSize=20000 ; Size in bytes of the TCP/IP receive buffer for each
 ; channel instance. Default is 32768.
Connect_Timeout=10000 ; Number of seconds before an attempt to connect the
 ; channel instance times out. Default is zero (no timeout).

QMErrorLog:
ErrorLogSize = 262144
ExcludeMessage = 7234
SuppressMessage = 9001,9002,9202
SuppressInterval = 30

```

*Figure 85. Example queue manager configuration file*

**Note:**

1. IBM MQ on the node uses the default locations for queue managers and the journals.
2. The queue manager saturn.queue.manager is the default queue manager for the node. The directory for files associated with this queue manager has been automatically transformed into a valid file name for the file system.
3. Because the IBM MQ configuration file is used to locate the data associated with queue managers, a nonexistent or incorrect configuration file can cause some or all IBM MQ commands to fail. Also, applications cannot connect to a queue manager that is not defined in the IBM MQ configuration file.

## Attributes for changing IBM MQ configuration information

On IBM MQ for Windows systems and on IBM MQ for Linux (x86 and x86-64 platforms) systems, modify configuration information using the IBM MQ Explorer. On other systems, modify the information by editing the `mqs.ini` configuration file.

See the following subtopics for attributes for specific components:

### Related concepts:

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

### Related reference:

“Changing queue manager configuration information” on page 782

The attributes described here modify the configuration of an individual queue manager. They override any settings for IBM MQ.

### Related information:

Planning

Administering IBM MQ

## All queue managers

Use the General and Extended IBM MQ properties page from the IBM MQ Explorer, or the `AllQueueManagers` stanza in the `mqs.ini` file to specify the following information about all queue managers.

### **DefaultPrefix=** *directory\_name*

This attribute specifies the path to the `qmgrs` directory, within which the queue manager data is kept.

If you change the default prefix for the queue manager, replicate the directory structure that was created at installation time.

In particular, you must create the `qmgrs` structure. Stop IBM MQ before changing the default prefix, and restart IBM MQ only after you have moved the structures to the new location and changed the default prefix.

**Note:** Do not delete the `/var/mqm/errors` directory on UNIX and Linux systems, or the `\errors` directory on Windows systems.

As an alternative to changing the default prefix, you can use the environment variable `MQSPREFIX` to override the `DefaultPrefix` for the `crtmqm` command.

Because of operating system restrictions, keep the supplied path sufficiently short so that the sum of the path length and any queue manager name is a maximum of 70 characters long.

### **ConvEBCDICNewline=**`NL_TO_LF`|`TABLE`|`ISO`

EBCDIC code pages contain a newline (NL) character that is not supported by ASCII code pages (although some ISO variants of ASCII contain an equivalent).

Use the `ConvEBCDICNewline` attribute to specify how IBM MQ is to convert the EBCDIC NL character into ASCII format.

#### **NL\_TO\_LF**

Convert the EBCDIC NL character (X'15') to the ASCII line feed character, LF (X'0A'), for all EBCDIC to ASCII conversions.

`NL_TO_LF` is the default.

**TABLE**

Convert the EBCDIC NL character according to the conversion tables used on your platform for all EBCDIC to ASCII conversions.

The effect of this type of conversion might vary from platform to platform and from language to language; even on the same platform, the behavior might vary if you use different CCSIDs.

**ISO**

Convert:

- ISO CCSIDs using the TABLE method
- All other CCSIDs using the NL\_TO\_CF method

Possible ISO CCSIDs are shown in Table 121.

Table 121. List of possible ISO CCSIDs

| CCSID | Code Set  |
|-------|-----------|
| 819   | ISO8859-1 |
| 912   | ISO8859-2 |
| 915   | ISO8859-5 |
| 1089  | ISO8859-6 |
| 813   | ISO8859-7 |
| 916   | ISO8859-8 |
| 920   | ISO8859-9 |
| 1051  | roman8    |

If the ASCII CCSID is not an ISO subset, ConvEBCDICNewline defaults to NL\_TO\_LF.

**Default queue manager**

Use the General IBM MQ properties page from the IBM MQ Explorer, or the DefaultQueueManager stanza in the mq.ini file to specify the default queue manager.

**Name=** *default\_queue\_manager*

The default queue manager processes any commands for which a queue manager name is not explicitly specified. The DefaultQueueManager attribute is automatically updated if you create a new default queue manager. If you inadvertently create a new default queue manager and then want to revert to the original, alter the DefaultQueueManager attribute manually.

**Exit properties**

Use the Extended IBM MQ properties page from the IBM MQ Explorer, or the ExitProperties stanza in the mq.ini file to specify configuration options used by queue manager exit programs.

**CLWLMode=** **SAFE** | **FAST**

The cluster workload (CLWL) exit allows you to specify which cluster queue in the cluster to open in response to an MQI call (for example, **MQOPEN** , **MQPUT** ). The CLWL exit runs either in FAST mode or SAFE mode depending on the value you specify on the CLWLMode attribute. If you omit the CLWLMode attribute, the cluster workload exit runs in SAFE mode.

**SAFE**

Run the CLWL exit in a separate process from the queue manager. This is the default.

If a problem arises with the user-written CLWL exit when running in SAFE mode, the following happens:

- The CLWL server process (amqzlw0) fails.
- The queue manager restarts the CLWL server process.



- The error is reported to you in the error log. If an MQI call is in progress, you receive notification in the form of a return code.

The integrity of the queue manager is preserved.

**Note:** Running the CLWL exit in a separate process can affect performance.

#### **FAST**

Run the cluster exit inline in the queue manager process.

Specifying this option improves performance by avoiding the process switching costs associated with running in SAFE mode, but does so at the expense of queue manager integrity. You should only run the CLWL exit in FAST mode if you are convinced that there are **no** problems with your CLWL exit, and you are particularly concerned about performance.

If a problem arises when the CLWL exit is running in FAST mode, the queue manager will fail and you run the risk of the integrity of the queue manager being compromised.

### **Log defaults for IBM MQ**

Use the Default log settings IBM MQ properties page from the IBM MQ Explorer, or the LogDefaults stanza in the mqs.ini file to specify information about log defaults for all queue managers.

If the stanza does not exist then the MQ defaults will be used. The log attributes are used as default values when you create a queue manager, but can be overridden if you specify the log attributes on the **crtmqm** command. See **crtmqm** for details of this command.

Once a queue manager has been created, the log attributes for that queue manager are taken from the settings described in “Queue manager logs” on page 786.

The default prefix (specified in “All queue managers” on page 775 ) and log path specified for the particular queue manager (specified in “Queue manager logs” on page 786 ) allow the queue manager and its log to be on different physical drives. This is the recommended method, although by default they are on the same drive.

For information about calculating log sizes, see “Calculating the size of the log” on page 1136.

**Note:** The limits given in the following parameter list are limits set by IBM MQ. Operating system limits might reduce the maximum possible log size.

#### **LogPrimaryFiles= 3 | 2-254 ( Windows ) | 2-510 ( UNIX and Linux systems)**

The log files allocated when the queue manager is created.

The minimum number of primary log files you can have is 2 and the maximum is 254 on Windows, or 510 on UNIX and Linux systems. The default is 3.

The total number of primary and secondary log files must not exceed 255 on Windows, or 511 on UNIX and Linux systems, and must not be less than 3.

The value is examined when the queue manager is created or started. You can change it after the queue manager has been created. However, a change in the value is not effective until the queue manager is restarted, and the effect might not be immediate.

#### **LogSecondaryFiles= 2 | 1-253 ( Windows ) | 1-509 ( UNIX and Linux systems)**

The log files allocated when the primary files are exhausted.

The minimum number of secondary log files is 1 and the maximum is 253 on Windows, or 509 on UNIX and Linux systems. The default number is 2.

The total number of primary and secondary log files must not exceed 255 on Windows, or 511 on UNIX and Linux systems, and must not be less than 3.

The value is examined when the queue manager is started. You can change this value, but changes do not become effective until the queue manager is restarted, and even then the effect might not be immediate.

**LogFilePages=** *number*

The log data is held in a series of files called log files. The log file size is specified in units of 4 KB pages.

The default number of log file pages is 4096, giving a log file size of 16 MB.

On UNIX and Linux systems the minimum number of log file pages is 64, and on Windows the minimum number of log file pages is 32; in both cases the maximum number is 65 535.

**Note:** The size of the log files specified during queue manager creation cannot be changed for a queue manager.

**LogType=** CIRCULAR | LINEAR

The type of log to be used. The default is CIRCULAR.

**CIRCULAR**

Start restart recovery using the log to roll back transactions that were in progress when the system stopped.

See “Types of logging” on page 1131 for a fuller explanation of circular logging.

**LINEAR**

For both restart recovery and media or forward recovery (creating lost or damaged data by replaying the contents of the log).

See “Types of logging” on page 1131 for a fuller explanation of linear logging.

If you want to change the default, you can either edit the LogType attribute, or specify linear logging using the **crtmqm** command. You cannot change the logging method after a queue manager has been created.

**LogBufferPages=** 0 | *0-4096*

The amount of memory allocated to buffer records for writing, specifying the size of the buffers in units of 4 KB pages.

The minimum number of buffer pages is 18 and the maximum is 4096. Larger buffers lead to higher throughput, especially for larger messages.

If you specify 0 (the default), the queue manager selects the size. In IBM WebSphere MQ Version 7.1 this is 512 (2048 KB).

If you specify a number in the range 1 through 17, the queue manager defaults to 18 (72 KB). If you specify a number in the range 18 and through 4096, the queue manager uses the number specified to set the memory allocated.

**LogDefaultPath=** *directory\_name*

The directory in which the log files for a queue manager reside. The directory resides on a local device to which the queue manager can write and, preferably, on a different drive from the message queues. Specifying a different drive gives added protection in case of system failure.

The default is:

- <DefaultPrefix>\log for IBM MQ for Windows where <DefaultPrefix> is the value specified on the DefaultPrefix attribute on the All Queue Managers IBM MQ properties page. This value is set at install time.
- /var/mqm/log for IBM MQ for UNIX and Linux systems

Alternatively, you can specify the name of a directory on the **crtmqm** command using the **-ld** flag. When a queue manager is created, a directory is also created under the queue manager directory, and

this is used to hold the log files. The name of this directory is based on the queue manager name. This ensures that the log file path is unique, and also that it conforms to any limitations on directory name lengths.

If you do not specify `-ld` on the `crtmqm` command, the value of the `LogDefaultPath` attribute in the `mqs.ini` file is used.

The queue manager name is appended to the directory name to ensure that multiple queue managers use different log directories.

When the queue manager is created, a `LogPath` value is created in the log attributes in the configuration information, giving the complete directory name for the queue manager's log. This value is used to locate the log when the queue manager is started or deleted.

#### **LogWriteIntegrity=SingleWrite|DoubleWrite| TripleWrite**

The method the logger uses to reliably write log records.

##### **TripleWrite**

This is the default method.

Note, that you can select **DoubleWrite**, but if you do so, the system interprets this as **TripleWrite**.

##### **SingleWrite**

You should use **SingleWrite**, only if the file-system and device hosting the IBM MQ recovery log explicitly guarantees the atomicity of 4KB writes.

That is, when a write of a 4KB page fails for any reason, the only two possible states are either the before image, or the after image. No intermediate state should be possible.

**Note:** If there is sufficient concurrency in your persistent workload, there is minimal potential benefit in setting anything other than the default value, **TripleWrite**.

### **Advanced Configuration and Power Interface (ACPI)**

Use the ACPI IBM MQ properties page from the IBM MQ Explorer, to specify how IBM MQ is to behave when the system receives a suspend request.

Windows supports the Advanced Configuration and Power Interface (ACPI) standard. This enables Windows users with ACPI enabled hardware to stop and restart channels when the system enters and resumes from suspend mode.

Note that the settings specified in the ACPI IBM MQ properties page are applied only when the Alert Monitor is running. The Alert Monitor icon is present on the taskbar if the Alert Monitor is running.

#### **DoDialog= Y | N**

Displays the dialog at the time of a suspend request.

#### **DenySuspend=Y | N**

Denies the suspend request. This is used if `DoDialog=N`, or if `DoDialog=Y` and a dialog cannot be displayed, for example, because your notebook lid is closed.

#### **CheckChannelsRunning=Y | N**

Checks whether any channels are running. The outcome can determine the outcome of the other settings.

The following table outlines the effect of each combination of these parameters:

| DoDialog | DenySuspend | CheckChannels Running | Action                                                                                                                 |
|----------|-------------|-----------------------|------------------------------------------------------------------------------------------------------------------------|
| N        | N           | N                     | Accept the suspend request.                                                                                            |
| N        | N           | Y                     | Accept the suspend request.                                                                                            |
| N        | Y           | N                     | Deny the suspend request.                                                                                              |
| N        | Y           | Y                     | If any channels are running deny the suspend request; if not accept the request.                                       |
| Y        | N           | N                     | Display the dialog (see Note ; accept the suspend request). This is the default.                                       |
| Y        | N           | Y                     | If no channels are running accept the suspend request; if they are display the dialog (see Note ; accept the request). |
| Y        | Y           | N                     | Display the dialog ( Note ; deny the suspend request).                                                                 |
| Y        | Y           | Y                     | If no channels are running accept the suspend request; if they are display the dialog ( Note ; deny the request).      |

**Note:** In cases where the action is to display the dialog, if the dialog cannot be displayed (for example because your notebook lid is closed), the DenySuspend option is used to determine whether the suspend request is accepted or denied.

## API exits

Use the MQ Explorer or the **amqmdain** command to change the entries for API exits.

Use the Exits IBM MQ properties page from the MQ Explorer, or the `ApiExitTemplate` and `ApiExitCommon` stanza in the `mqs.ini` file to identify API exit routines for all queue managers. On Windows systems, you can also use the **amqmdain** command to change the entries for API exits. (To identify API exit routines for individual queue managers, you use the `ApiExitLocal` stanza, as described in “API exits” on page 796.)

For a complete description of the attributes for these stanzas, see Configuring API exits.

## Queue managers

There is one `QueueManager` stanza for every queue manager. Use the stanza to specify the location of the queue manager directory.

On Windows, UNIX and Linux systems, there is one `QueueManager` stanza for every queue manager. These attributes specify the queue manager name, and the name of the directory containing the files associated with that queue manager. The name of the directory is based on the queue manager name, but is transformed if the queue manager name is not a valid file name. See, Understanding IBM MQ file names for more information about name transformation.

**Name=** *queue\_manager\_name*

The name of the queue manager.

**Prefix=** *prefix*

Where the queue manager files are stored. By default, this value is the same as the value specified on the `DefaultPrefix` attribute of the All Queue Managers information.

**Directory=** *name*

The name of the subdirectory under the `<prefix>\QMGRS` directory where the queue manager files are stored. This name is based on the queue manager name, but can be transformed if there is a duplicate name or if the queue manager name is not a valid file name.

**DataPath= path**

An explicit data path provided when the queue manager was created, this overrides Prefix and Directory as the path to the queue manager data.

**InstallationName= name**

The name of the IBM MQ installation associated with this queue manager. Commands from this installation must be used when interacting with this queue manager. If no InstallationName value is present, the queue manager is associated with an installation of IBM MQ earlier than Version 7.1.

**Related concepts:**

“Associating a queue manager with an installation” on page 1034

When you create a queue manager, it is automatically associated with the installation that issued the **crtmqm** command. On UNIX, Linux, and Windows, you can change the installation associated with a queue manager using the **setmqm** command.

**Security**

Use the Security stanza in the `qm.ini` file to specify options for the Object Authority Manager (OAM).

**ClusterQueueAccessControl=RQMName|Xmitq**

Set this attribute to check the access control of cluster queues or fully qualified queues hosted on cluster queue managers.

**RQMName**

The profiles checked for access control of remotely hosted queues are named queues or named queue manager profiles.

**Xmitq**

The profiles checked for access control of remotely hosted queues are resolved to the `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

`Xmitq` is the default value.

**GroupModel=GlobalGroups**

This attribute determines whether the OAM checks global groups when determining the group membership of a user on Windows.

The default is not to check global groups.

**GlobalGroups**

The OAM checks global groups.

With `GlobalGroups` set, the authorization commands, **setmqaut**, **dspmqaut**, and **dmpmqaut** accept global groups names; see the **setmqaut -g** parameter.

**Note:** Setting the `ClusterQueueAccessControl=RQMName` and having a custom implementation of the Authorization Service at less than `MQZAS_VERSION_6` results in the queue manager not starting. In this instance, either set `ClusterQueueAccessControl=Xmitq` or upgrade the custom Authorization Service to `MQZAS_VERSION_6` or greater.

## Changing queue manager configuration information

The attributes described here modify the configuration of an individual queue manager. They override any settings for IBM MQ.

On UNIX and Linux systems, you modify queue manager configuration information by editing the `qm.ini` configuration file. When you are defining a stanza in `qm.ini`, you do not need to start each item on a new line. You can use either a semicolon (;) or a hash character (#) to indicate a comment.

On Windows and Linux x86-64 systems, you can modify some configuration information by using the MQ Explorer. However, because there are significant implications to changing installable services and their components, the installable services are read-only in the MQ Explorer. You must therefore make any changes to installable services by using **regedit** on Windows, and by editing the `qm.ini` file on UNIX and Linux.

For more details on changing queue manager configuration information, see the following subtopics:

### Related concepts:

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

### Related reference:

“Attributes for changing IBM MQ configuration information” on page 775

On IBM MQ for Windows systems and on IBM MQ for Linux (x86 and x86-64 platforms) systems, modify configuration information using the IBM MQ Explorer. On other systems, modify the information by editing the `mqs.ini` configuration file.

### Related information:

Planning

Administering IBM MQ

## Access Mode

**Access Mode** applies to Windows servers only. The `AccessMode` stanza is set by the `-a [r]` option on the `crtmqm` command. Do not change the `AccessMode` stanza after the queue manager has been created.

Use the access group ( `-a [r]` ) option of the `crtmqm` command to specify a Windows security group, members of which will be granted full access to all queue manager data files. The group can either be a local or global group, depending upon the syntax used. Valid syntax for the group name is as follows:

*LocalGroup*

*Domain name\GlobalGroup name*

*GlobalGroup name @ Domain name*

You must define the additional access group before running the `crtmqm` command with the `-a [r]` option.

If you specify the group using `-ar` instead of `-a`, the local `mqm` group is not granted access to the queue manager data files. Use this option, if the file system hosting the queue manager data files does not support access control entries for locally defined groups.

The group is typically a global security group, which is used to provide multi-instance queue managers with access to a shared queue manager data and logs folder. Use the additional security access group to set read and write permissions on the folder or to share containing queue manager data and log files.

The additional security access group is an alternative to using the local group named `mqm` to set permissions on the folder containing queue manager data and logs. Unlike the local group `mqm`, you can

make the additional security access group a local or a global group. It must be a global group to set permissions on the shared folders that contain the data and log files used by multi-instance queue managers.

The Windows operating system checks the access permissions to read and write queue manager data and log files. It checks the permissions of the user ID that is running queue manager processes. The user ID that is checked depends on whether you started the queue manager as a service or you started it interactively. If you started the queue manager as a service, the user ID checked by the Windows system is the user ID you configured with the Prepare IBM MQ wizard. If you started the queue manager interactively, the user ID checked by the Windows system is the user ID that ran the **strmqm** command.

The user ID must be a member of the local **mqm** group to start the queue manager. If the user ID is a member of the additional security access group, the queue manager can read and write files that are given permissions by using the group.

**Restriction:** You can specify an additional security access group only on Windows operating system. If you specify an additional security access group on other operating systems, the **crtmqm** command returns an error.

**Related concepts:**

“Secure unshared queue manager data and log directories and files on Windows” on page 1106  
This topic describes how you can secure an alternative location for queue manager data and log files, both by using the local **mqm** group and an alternative security group.

“Secure shared queue manager data and log directories and files on Windows” on page 1103  
This topic describes how you can secure a shared location for queue manager data and log files using a global alternative security group. You can share the location between different instances of a queue manager running on different servers.

**Related tasks:**

“Create a multi-instance queue manager on domain workstations or servers” on page 1080  
An example shows how to set up a multi-instance queue manager on Windows on a workstation or a server that is part of a Windows domain. The server does not have to be a domain controller. The setup demonstrates the concepts involved, rather than being production scale. The example is based on Windows Server 2008. The steps might differ on other versions of Windows Server.

**Related information:**

`crtmqm`

## Installable services

You change installable services on Windows by using **regedit**, and on UNIX and Linux by using the Service stanza in the `qm.ini` file.

**Note:** There are significant implications to changing installable services and their components. For this reason, the installable services are read-only in the MQ Explorer.

To change installable services on Windows systems, use **regedit**, or on UNIX and Linux systems, use the Service stanza in the `qm.ini` file. For each component within a service, you must also specify the name and path of the module containing the code for that component. On UNIX and Linux systems, use the `ServiceComponent` stanza for this.

**Name= AuthorizationService |NameService**

The name of the required service.

**AuthorizationService**

For IBM MQ, the Authorization Service component is known as the object authority manager, or OAM. The `AuthorizationService` stanza and its associated `ServiceComponent` stanza are added automatically when the queue manager is created. Add other `ServiceComponent` stanzas manually.

### **NameService**

No name service is provided by default. If you require a name service, you must add the NameService stanza manually.

### **EntryPoints=** *number-of-entries*

The number of entry points defined for the service.

This includes the initialization and termination entry points.

Windows

### **SecurityPolicy= Default | NTSIDsRequired**

On Windows systems, the SecurityPolicy attribute applies only if the service specified is the default authorization service, that is, the OAM. The SecurityPolicy attribute allows you to specify the security policy for each queue manager.

The possible values are:

#### **Default**

Use the default security policy to take effect. If a Windows security identifier (NT SID) is not passed to the OAM for a particular user ID, an attempt is made to obtain the appropriate SID by searching the relevant security databases.

#### **NTSIDsRequired**

Pass an NT SID to the OAM when performing security checks.

See Windows security identifiers (SIDs) for more information.

See also Configuring authorization service stanzas: Windows systems.

Linux

UNIX

### **SecurityPolicy=user|group| default**

On UNIX and Linux systems the value specifies whether the queue manager uses user-based or group-based authorization. Values are not case sensitive.

If you do not include this attribute, default is used, which uses group-based authorization. Restart the queue manager for changes to become effective. See also Configuring authorization service stanzas: UNIX and Linux systems.

### **SharedBindingsUserId=** *user-type*

The SharedBindingsUserId attribute applies only if the service specified is the default authorization service, that is, the OAM. The SharedBindingsUserId attribute is used with relation to shared bindings only. This value allows you to specify whether the *UserIdentifier* field in the *IdentityContext* structure, from the MQZ\_AUTHENTICATE\_USER function, is the effective user Id or the real user Id.

For information on the MQZ\_AUTHENTICATE\_USER function, see MQZ\_AUTHENTICATE\_USER - Authenticate user.

The possible values are:

#### **Default**

The value of the *UserIdentifier* field is set as the real user Id.

#### **Real**

The value of the *UserIdentifier* field is set as the real user Id.

#### **Effective**

The value of the *UserIdentifier* field is set as the effective user Id.

### **FastpathBindingsUserId=** *user-type*

The FastpathBindingsUserId attribute applies only if the service specified is the default authorization service, that is, the OAM. The FastpathBindingsUserId attribute is used with relation to fastpath bindings only. This value allows you to specify whether the *UserIdentifier* field in the *IdentityContext* structure, from the MQZ\_AUTHENTICATE\_USER function, is the effective user Id or the real user Id.

For information on the MQZ\_AUTHENTICATE\_USER function, see MQZ\_AUTHENTICATE\_USER - Authenticate user.



The possible values are:

**Default**

The value of the *UserIdentifier* field is set as the real user ID.

**Real**

The value of the *UserIdentifier* field is set as the real user ID.

**Effective**

The value of the *UserIdentifier* field is set as the effective user ID.

**IsolatedBindingsUserId= *user-type***

The **IsolatedBindingsUserId** attribute applies only if the service specified is the default authorization service, that is, the OAM. The **IsolatedBindingsUserId** attribute is used with relation to isolated bindings only. This value allows you to specify whether the *UserIdentifier* field in the *IdentityContext* structure, from the MQZ\_AUTHENTICATE\_USER function, is the effective user Id or the real user Id.

For information on the MQZ\_AUTHENTICATE\_USER function, see MQZ\_AUTHENTICATE\_USER - Authenticate user.

The possible values are:

**Default**

The value of the *UserIdentifier* field is set as the effective user Id.

**Real**

The value of the *UserIdentifier* field is set as the real user Id.

**Effective**

The value of the *UserIdentifier* field is set as the effective user Id.

For more information about installable services and components, see Installable services and components for UNIX, Linux and Windows .

For more information about security services in general, see Setting up security on UNIX and Linux systems.

**Related information:**

Installable services reference information

**Service components:**

You must specify service component information when you add a new installable service. On Windows systems use **regedit** , and on UNIX and Linux systems use the ServiceComponent stanza in the qm.ini file. The authorization service stanza is present by default, and the associated component, the OAM, is active.

Specify the service components as follows:

**Service= *service\_name***

The name of the required service. This must match the value specified on the Name attribute of the Service configuration information.

**Name= *component\_name***

The descriptive name of the service component. This must be unique and contain only characters that are valid for the names of IBM MQ objects (for example, queue names). This name occurs in operator messages generated by the service. We recommend that this name begins with a company trademark or similar distinguishing string.

**Module= *module\_name***

The name of the module to contain the code for this component. This must be a full path name.

**ComponentDataSize=** *size*

The size, in bytes, of the component data area passed to the component on each call. Specify zero if no component data is required.

For more information about installable services and components, see *Installable services and components for UNIX, Linux and Windows* .

## Queue manager logs

Use the Log queue manager properties page from the IBM MQ Explorer, or the Log stanza in the qm.ini file, to specify information about logging on a queue manager.

By default, these settings are inherited from the settings specified for the default log settings for the queue manager (described in “Log defaults for IBM MQ” on page 777 ). Change these settings only if you want to configure this queue manager in a different way.

For information about calculating log sizes, see “Calculating the size of the log” on page 1136.

**Note:** The limits given in the following parameter list are set by IBM MQ. Operating system limits might reduce the maximum possible log size.

**LogPrimaryFiles=** 3 | 2-254 ( Windows ) | 2-510 ( UNIX and Linux systems)

The log files allocated when the queue manager is created.

The minimum number of primary log files you can have is 2 and the maximum is 254 on Windows, or 510 on UNIX and Linux systems. The default is 3.

The total number of primary and secondary log files must not exceed 255 on Windows, or 511 on UNIX and Linux systems, and must not be less than 3.

The value is examined when the queue manager is created or started. You can change it after the queue manager has been created. However, a change in the value is not effective until the queue manager is restarted, and the effect might not be immediate.

**LogSecondaryFiles=** 2 | 1-253 ( Windows ) | 1-509 ( UNIX and Linux systems)

The log files allocated when the primary files are exhausted.

The minimum number of secondary log files is 1 and the maximum is 253 on Windows, or 509 on UNIX and Linux systems. The default number is 2.

The total number of primary and secondary log files must not exceed 255 on Windows, or 511 on UNIX and Linux systems, and must not be less than 3.

The value is examined when the queue manager is started. You can change this value, but changes do not become effective until the queue manager is restarted, and even then the effect might not be immediate.

**LogFilePages=** *number*

The log data is held in a series of files called log files. The log file size is specified in units of 4 KB pages.

The default number of log file pages is 4096, giving a log file size of 16 MB.

On UNIX and Linux systems the minimum number of log file pages is 64, and on Windows the minimum number of log file pages is 32; in both cases the maximum number is 65 535.

**Note:** The size of the log files specified during queue manager creation cannot be changed for a queue manager.

**LogType=** CIRCULAR | LINEAR

The type of logging to be used by the queue manager. You cannot change the type of logging to be

used once the queue manager has been created. Refer to the description of the LogType attribute in “Log defaults for IBM MQ” on page 777 for information about creating a queue manager with the type of logging you require.

### **CIRCULAR**

Start restart recovery using the log to roll back transactions that were in progress when the system stopped.

See “Types of logging” on page 1131 for a fuller explanation of circular logging.

### **LINEAR**

For both restart recovery and media or forward recovery (creating lost or damaged data by replaying the contents of the log).

See “Types of logging” on page 1131 for a fuller explanation of linear logging.

### **LogBufferPages= 0 | 0-4096**

The amount of memory allocated to buffer records for writing, specifying the size of the buffers in units of 4 KB pages.

The minimum number of buffer pages is 18 and the maximum is 4096. Larger buffers lead to higher throughput, especially for larger messages.

If you specify 0 (the default), the queue manager selects the size. In IBM WebSphere MQ Version 7.1 this is 512 (2048 KB).

If you specify a number in the range 1 through 17, the queue manager defaults to 18 (72 KB). If you specify a number in the range 18 through 4096, the queue manager uses the number specified to set the memory allocated.

The value is examined when the queue manager is started. The value can be increased or decreased within the limits stated. However, a change in the value is not effective until the next time the queue manager is started.

### **LogPath= *directory\_name***

The directory in which the log files for a queue manager reside. This must exist on a local device to which the queue manager can write and, preferably, on a different drive from the message queues. Specifying a different drive gives added protection in case of system failure.

The default is:

- C:\ProgramData\IBM\MQ\log in IBM MQ for Windows.
- /var/mqm/log in IBM MQ for UNIX and Linux systems.

You can specify the name of a directory on the **crtmqm** command using the **-ld** flag. When a queue manager is created, a directory is also created under the queue manager directory, and this is used to hold the log files. The name of this directory is based on the queue manager name. This ensures that the log file path is unique, and also that it conforms to any limitations on directory name lengths.

If you do not specify **-ld** on the **crtmqm** command, the value of the LogDefaultPath attribute is used.

In IBM MQ for UNIX and Linux systems, user ID **mqm** and group **mqm** must have full authorities to the log files. If you change the locations of these files, you must give these authorities yourself. This is not required if the log files are in the default locations supplied with the product.

### **LogWriteIntegrity=SingleWrite|DoubleWrite| TripleWrite**

The method the logger uses to reliably write log records.

### **TripleWrite**

This is the default method.

Note, that you can select **DoubleWrite**, but if you do so, the system interprets this as **TripleWrite**.

## SingleWrite

You should use **SingleWrite**, only if the file-system and device hosting the IBM MQ recovery log explicitly guarantees the atomicity of 4KB writes.

That is, when a write of a 4KB page fails for any reason, the only two possible states are either the before image, or the after image. No intermediate state should be possible.

**Note:** If there is sufficient concurrency in your persistent workload, there is minimal potential benefit in setting anything other than the default value, **TripleWrite**.

## Restricted mode

This option applies to UNIX and Linux systems only. The RestrictedMode stanza is set by the -g option on the **crtmqm** command. Do not change this stanza after the queue manager has been created. If you do not use the -g option, the stanza is not created in the qm.ini file.

There are some directories under which IBM MQ applications create files while they are connected to the queue manager within the queue manager data directory. In order for applications to create files in these directories, they are granted world write access:

- /var/mqm/sockets/QMgrName/@ipcc/ssem/hostname/
- /var/mqm/sockets/QMgrName/@app/ssem/hostname/
- /var/mqm/sockets/QMgrName/zsocketapp/hostname/

where <QMGRNAME> is the name of the queue manager, and <hostname> is the host name.

On some systems, it is unacceptable to grant all users write access to these directories. For example, those users who do not need access the queue manager. Restricted mode modifies the permissions of the directories that store queue manager data. The directories can then only be accessed by members of the specified application group. The permissions on the System V IPC shared memory used to communicate with the queue manager are also modified in the same way.

The application group is the name of the group with members that have permission to do the following things:




- run MQI applications
- update all IPCC resources
- change the contents of some queue manager directories

To use restricted mode for a queue manager:

- The creator of the queue manager must be in the mqm group and in the application group.
- The mqm user ID must be in the application group.
- All users who want to administer the queue manager must be in the mqm group and in the application group.
- all users who want to run IBM MQ applications must be in the application group.

Any MQCONN or MQCONNX call issued by a user who is not in the application group failed with reason code MQRC\_Q\_MGR\_NOT\_AVAILABLE.

Restricted mode operates with the IBM MQ authorization service. Therefore you must also grant users the authority to connect to IBM MQ and access the resources they require using the IBM MQ authorization service.

   Further information about configuring the IBM MQ authorization service can be found in Setting up security on Windows, UNIX and Linux systems.

Only use IBM MQ restricted mode when the control provided by the authorization service does not provide sufficient isolation of queue manager resources.

## XA resource managers

Use the XA resource manager queue manager properties page from the IBM MQ Explorer, or the XAResourceManager stanza in the qm.ini file, to specify the following information about the resource managers involved in global units of work coordinated by the queue manager.

Add XA resource manager configuration information manually for each instance of a resource manager participating in global units of work; no default values are supplied.

See Database coordination for more information about resource manager attributes.

### **Name= *name* (mandatory)**

This attribute identifies the resource manager instance.

The Name value can be up to 31 characters in length. You can use the name of the resource manager as defined in its XA-switch structure. However, if you are using more than one instance of the same resource manager, you must construct a unique name for each instance. You can ensure uniqueness by including the name of the database in the Name string, for example.

IBM MQ uses the Name value in messages and in output from the **dspmqtrn** command.

Do not change the name of a resource manager instance, or delete its entry from the configuration information, once the associated queue manager has started and the resource manager name is in effect.

### **SwitchFile= *name* (mandatory)**

The fully-qualified name of the load file containing the resource manager's XA switch structure.

If you are using a 64-bit queue manager with 32-bit applications, the name value should contain only the base name of the load file containing the resource manager's XA switch structure.

The 32-bit file will be loaded into the application from the path specified by **ExitsDefaultPath** .

The 64-bit file will be loaded into the queue manager from the path specified by **ExitsDefaultPath64** .

### **XAOpenString= *string* (optional)**

The string of data to be passed to the resource manager's xa\_open entry point. The contents of the string depend on the resource manager itself. For example, the string could identify the database that this instance of the resource manager is to access. For more information about defining this attribute, see:

- Adding resource manager configuration information for Db2
- Adding resource manager configuration information for Oracle
- Adding resource manager configuration information for Sybase
- Adding resource manager configuration information for Informix®

and consult your resource manager documentation for the appropriate string.

### **XACloseString= *string* (optional)**

The string of data to be passed to the resource manager's xa\_close entry point. The contents of the string depend on the resource manager itself. For more information about defining this attribute, see:

- Adding resource manager configuration information for Db2
- Adding resource manager configuration information for Oracle
- Adding resource manager configuration information for Sybase
- Adding resource manager configuration information for Informix

and consult your database documentation for the appropriate string.

**ThreadOfControl=THREAD | PROCESS**

This attribute is mandatory for IBM MQ for Windows. The queue manager uses this value for serialization when it needs to call the resource manager from one of its own multithreaded processes.

**THREAD**

The resource manager is fully *thread aware*. In a multithreaded IBM MQ process, XA function calls can be made to the external resource manager from multiple threads at the same time.

**PROCESS**

The resource manager is not *thread safe*. In a multithreaded IBM MQ process, only one XA function call at a time can be made to the resource manager.

The ThreadOfControl entry does not apply to XA function calls issued by the queue manager in a multithreaded application process. In general, an application that has concurrent units of work on different threads requires this mode of operation to be supported by each of the resource managers.

**Attributes of the channels stanza**

These attributes determine the configuration of a channel.

This information is not applicable to IBM MQ for the z/OS platform.

Use the Channels queue manager properties page from the IBM MQ Explorer, or the CHANNELS stanza in the `qm.ini` file, to specify information about channels.

**MaxChannels= 100 | number**

The maximum number of *current* channels allowed.

The value must be in the range 1 - 65535. The default is 100.

**MaxActiveChannels= MaxChannels\_value**

The maximum number of channels allowed to be *active* at any time. The default is the value specified for the MaxChannels attribute.

**MaxInitiators= 3 | number**

The maximum number of initiators. The default and maximum value is 3.

**MQIBindType=FASTPATH | STANDARD**

The binding for applications:

**FASTPATH**

Channels connect using **MQCONN** FASTPATH; there is no agent process.

**STANDARD**

Channels connect using STANDARD.

**PipeLineLength= 1 | number**

The maximum number of concurrent threads a channel will use. The default is 1. Any value greater than 1 is treated as 2.

When you use pipelining, configure the queue managers at both ends of the channel to have a *PipeLineLength* greater than 1.

**Note:** Pipelining is only effective for TCP/IP channels.

**AdoptNewMCA= NO | SVR | SDR | RCVR | CLUSRCVR | ALL | FASTPATH**

If IBM MQ receives a request to start a channel, but finds that an instance of the channel is already running, in some cases the existing channel instance must be stopped before the new one can start. The AdoptNewMCA attribute allows you to control which types of channels can be ended in this way.

If you specify the AdoptNewMCA attribute for a particular channel type, but the new channel fails to start because a matching channel instance is already running:

1. The new channel tries to stop the previous one by requesting it to end.

2. If the previous channel server does not respond to this request by the time the `AdoptNewMCATimeout` wait interval expires, the thread or process for the previous channel server is ended.
3. If the previous channel server has not ended after step 2, and after the `AdoptNewMCATimeout` wait interval expires for a second time, IBM MQ ends the channel with a `CHANNEL IN USE` error.

The `AdoptNewMCA` functionality applies to server, sender, receiver, and cluster-receiver channels. In the case of a sender or server channel, only one instance of a channel with a particular name can be running in the receiving queue manager. In the case of a receiver or cluster-receiver channel, multiple instances of a channel with a particular name might be running in the receiving queue manager, but only one instance can run at any one time from a particular remote queue manager.

**Note:** `AdoptNewMCA` is not supported on requester or server-connection channels.

Specify one or more values, separated by commas or blanks, from the following list:

**NO** The `AdoptNewMCA` feature is not required. This is the default.

**SVR**

Adopt server channels.

**SDR**

Adopt sender channels.

**RCVR**

Adopt receiver channels.

**CLUSRCVR**

Adopt cluster receiver channels.

**ALL**

Adopt all channel types except `FASTPATH` channels.

**FASTPATH**

Adopt the channel if it is a `FASTPATH` channel. This happens only if the appropriate channel type is also specified, for example: `AdoptNewMCA=RCVR,SVR,FASTPATH`.

**Attention!:** The `AdoptNewMCA` attribute might behave in an unpredictable fashion with `FASTPATH` channels. Exercise great caution when enabling the `AdoptNewMCA` attribute for `FASTPATH` channels.

**`AdoptNewMCATimeout= 60 | 1 - 3600`**

The amount of time, in seconds, that the new channel instance waits for the old channel instance to end. Specify a value in the range 1 - 3600. The default value is 60.

**`AdoptNewMCACheck=QM|ADDRESS|NAME|ALL`**

The type of checking required when enabling the `AdoptNewMCA` attribute. If possible, perform full checking to protect your channels from being shut down, inadvertently or maliciously. At the very least, check that the channel names match.

Specify one or more of the following values, separated by commas or blanks in the case of `QM`, `NAME`, or `ALL`:

**QM** Check that the queue manager names match.

Note that the queue manager name itself is matched, not the `QMID`.

**ADDRESS**

Check the communications source IP address. For example, the TCP/IP address.

**Note:** Comma separated `CONNNAME` values apply to target addresses and are, therefore, not relevant to this option.

In the case that a multi-instance queue manager fails over from `hosta` to `hostb`, any outbound channels from that queue manager will use the source IP address of `hostb`. If this is different from `hosta`, then `AdoptNewMCACheck=ADDRESS` fails to match.

You can use SSL or TLS with mutual authentication to prevent an attacker from disrupting an existing running channel. Alternatively, use an HACMP type solution with IP-takeover instead of multi-instance queue managers, or use a network load balancer to mask the source IP address.

#### NAME

Check that the channel names match.

#### ALL

Check for matching queue manager names, the communications address, and for matching channel names.

The default is `AdoptNewMCACheck=NAME,ADDRESS,QM`.

#### **V8.0.0.5** `Ch1authEarlyAdopt=Y|N|E`

When you use the **ADOPTCTX(YES)** parameter on an authentication information object, the security context is set as the user ID that is presented in the MQCSP structure, when validated by a password. In this case, another security context cannot be adopted, unless you set the **Ch1authEarlyAdopt** parameter.

Valid values for **Ch1authEarlyAdopt** are the following values:

**Y** The channel validates and adopts user ID and password credentials that have been provided by an application using queue manager connection authentication before applying channel authentication rules. In this mode of operation, channel authentication rules match against the user ID resulting from connection authentication checks.

**N** The channel delays connection authentication validation of user ID and password credentials that have been provided by an application until after channel authentication rules have been applied. Note that in this mode of operation, channel authentication blocking and mapping rules cannot consider the results of user ID and password validation.

**E** **V8.0.0.5** From IBM MQ Version 8.0.0, Fix Pack 5, when security exits are enabled for a channel, allow the adoption of another security context when you use the **ADOPTCTX(YES)** parameter in an authentication information object. If you use this value when security exits are not in use, this value is the same as **Y**.

**V8.0.0.7** From IBM MQ Version 8.0.0, Fix Pack 7, the behavior is the same as the value **Y**.

For example, the default authentication information object is set to **ADOPTCTX(YES)**, and the user `fred` is logged in. The following two `CHLAUTH` rules are configured:

```
SET CHLAUTH('MY.CHLAUTH') TYPE(ADDRESSMAP) DESCR('Block all access by
default') ADDRESS('*') USERSRC(NOACCESS) ACTION(REPLACE)
SET CHLAUTH('MY.CHLAUTH') TYPE(USERMAP) DESCR('Allow user bob and force
CONNAUTH') CLNTUSER('bob') CHCKCLNT(REQUIRED) USERSRC(CHANNEL)
```

The following command is issued, with the intention of authenticating the command as the adopted security context of the user `bob`:

```
runmqsc -c -u bob QMGR
```

In fact, the queue manager uses the security context of `fred`, not `bob`, and the connection fails.

To use the security context of `bob`, **Ch1authEarlyAdopt** must be set to **Y**.

#### **PasswordProtection = Compatible | always | optional**

From IBM MQ Version 8.0, set protected passwords in the MQCSP structure, rather than using SSL or TLS.



MQCSP password protection is useful for test and development purposes as using MQCSP password protection is simpler than setting up SSL/TLS encryption, but not as secure.

For more information, see MQCSP password protection.

**Related concepts:**

“Channel states” on page 834

A channel can be in one of many states at any time. Some states also have substates. From a given state a channel can move into other states.

## TCP, LU62, and NETBIOS

Use these queue manager properties pages, or stanzas in the qm.ini file, to specify network protocol configuration parameters. They override the default attributes for channels.

### TCP

Use the TCP queue manager properties page from the IBM MQ Explorer, or the TCP stanza in the qm.ini file, to specify Transmission Control Protocol/Internet Protocol (TCP/IP) configuration parameters.

**Port= 1414** | *port\_number*

The default port number, in decimal notation, for TCP/IP sessions. The *well known* port number for IBM MQ is 1414.

**Library1= DLLName1** ( IBM MQ for Windows only)

The name of the TCP/IP sockets DLL.

The default is WSOCK32.

**KeepAlive= NO** | YES

Switch the KeepAlive function on or off. KeepAlive=YES causes TCP/IP to check periodically that the other end of the connection is still available. If it is not, the channel is closed.

**ListenerBacklog=number**

Override the default number of outstanding requests for the TCP/IP listener.

When receiving on TCP/IP, a maximum number of outstanding connection requests is set. This can be considered to be a backlog of requests waiting on the TCP/IP port for the listener to accept the request. The default listener backlog values are shown in Table 122.

Table 122. Default outstanding connection requests (TCP)

| Platform            | Default ListenerBacklog value |
|---------------------|-------------------------------|
| Windows Server      | 100                           |
| Windows Workstation | 5                             |
| Linux               | 100                           |
| Solaris             | 100                           |
| HP-UX               | 20                            |
| AIX V5.3 or later   | 100                           |

**Note:** Some operating systems support a larger value than the default shown. Use this to avoid reaching the connection limit.

Conversely, some operating systems might limit the size of the TCP backlog, so the effective TCP backlog could be smaller than requested here.

If the backlog reaches the values shown in Table 122, the TCP/IP connection is rejected and the channel cannot start. For message channels, this results in the channel going into a RETRY state

and retrying the connection at a later time. For client connections, the client receives an MQRC\_Q\_MGR\_NOT\_AVAILABLE reason code from MQCONN and retries the connection at a later time.

The following group of properties can be used to control the size of buffers used by TCP/IP. The values are passed directly to the TCP/IP layer of the operating system. Great care should be taken when using these properties. If the values are set incorrectly it can adversely affect the TCP/IP performance. For further information about how this affects performance refer to the TCP/IP documentation for your environment. A value of zero indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by IBM MQ.

**Connect\_Timeout= 0 | number**

The number of seconds before an attempt to connect the socket times out. The default value of zero specifies that there is no connect timeout.

IBM MQ channel processes connect over nonblocking sockets. Therefore, if the other end of the socket is not ready, connect() returns immediately with *EINPROGRESS* or *EWOULDBLOCK*. Following this, connect will be attempted again, up to a total of 20 such attempts, when a communications error is reported.

If **Connect\_Timeout** is set to a non-zero value, IBM MQ waits for the stipulated period over select() call for the socket to get ready. This increases the chances of success of a subsequent connect() call. This option might be beneficial in situations where connects would require some waiting period, due to high load on the network.

**SndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the sending end of channels. This stanza value can be overridden by a stanza more specific to the channel type, for example RcvSndBuffSize. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**RcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the receiving end of channels. This stanza value can be overridden by a stanza more specific to the channel type, for example RcvRcvBuffSize. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**RcvSndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the sender end of a receiver channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**RcvRcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the receiving end of a receiver channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**SvrSndBuffSize=number | 0**

The size in bytes of the TCP/IP send buffer used by the server end of a client-connection server-connection channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.

**SvrRcvBuffSize=number | 0**

The size in bytes of the TCP/IP receive buffer used by the server end of a client-connection server-connection channel. If the value is set as zero, the operating system defaults are used. If no value is set, then the IBM MQ default, 32768, is used.



**LU62 ( IBM MQ for Windows only)**

Use the LU6.2 queue manager properties page from the IBM MQ Explorer, or the LU62 stanza in the qm.ini file, to specify SNA LU 6.2 protocol configuration parameters.

**TPName**

The TP name to start on the remote site.

**Library1= *DLLName 1***

The name of the APPC DLL.

The default value is WCPIC32.

**Library2= *DLLName2***

The same as Library1, used if the code is stored in two separate libraries.

The default value is WCPIC32.

**Windows****NETBIOS ( IBM MQ for Windows only)**

Use the Netbios queue manager properties page from the IBM MQ Explorer, or the NETBIOS stanza in the qm.ini file, to specify NetBIOS protocol configuration parameters.

**LocalName= *name***

The name by which this machine is known on the LAN.

**AdapterNum= 0 | *adapter\_number***

The number of the LAN adapter. The default is adapter 0.

**NumSess= 1 | *number\_of\_sessions***

The number of sessions to allocate. The default is 1.

**NumCmds= 1 | *number\_of\_commands***

The number of commands to allocate. The default is 1.

**NumNames= 1 | *number\_of\_names***

The number of names to allocate. The default is 1.

**Library1= *DLLName1***

The name of the NetBIOS DLL.

The default value is NETAPI32.

**Windows****SPX ( IBM MQ for Windows only)**

Use the SPX queue manager properties page from the IBM MQ Explorer, or the SPX stanza in the qm.ini file, to specify SPX protocol configuration parameters.

**Socket= 5E86 | *socket\_number***

The SPX socket number in hexadecimal notation. The default is X'5E86'.

**BoardNum= 0 | *adapter\_number***

The LAN adapter number. The default is adapter 0.

**KeepAlive=NO|YES**

Switch the KeepAlive function on or off.

KeepAlive=YES causes SPX to check periodically that the other end of the connection is still available. If it is not, the channel is closed.

**Library1= *DLLName1***

The name of the SPX DLL.

The default is WSOCK32.DLL.

**Library2= *DLLName2***

The same as LibraryName1, used if the code is stored in two separate libraries.

The default is WSOCK32.DLL.

**ListenerBacklog=number**

Override the default number of outstanding requests for the SPX listener.

When receiving on SPX, a maximum number of outstanding connection requests is set. This can be considered to be a backlog of requests waiting on the SPX socket for the listener to accept the request. The default listener backlog values are shown in Table 123.

Table 123. Default outstanding connection requests (SPX)

| Platform            | Default ListenerBacklog value |
|---------------------|-------------------------------|
| Windows Server      | 100                           |
| Windows Workstation | 5                             |

**Note:** Some operating systems support a larger value than the default shown. Use this to avoid reaching the connection limit.

Conversely, some operating systems might limit the size of the SPX backlog, so the effective SPX backlog could be smaller than requested here.

If the backlog reaches the values shown in Table 123, the SPX connection is rejected and the channel cannot start. For message channels, this results in the channel going into a RETRY state and retrying the connection at a later time. For client connections, the client receives an MQRC\_Q\_MGR\_NOT\_AVAILABLE reason code from MQCONN and should retry the connection at a later time.

## Exit path

Use the Exits queue manager properties page from the IBM MQ Explorer, or the ExitPath stanza in the qm.ini file to specify the path for user exit programs on the queue manager system.

### ExitsDefaultPath= *string*

The ExitsDefaultPath attribute specifies the location of:

- 32-bit channel exits for clients
- 32-bit channel exits and data conversion exits for servers
- Unqualified XA switch load files

### ExitsDefaultPath64= *string*

The ExitsDefaultPath64 attribute specifies the location of:

- 64-bit channel exits for clients
- 64-bit channel exits and data conversion exits for servers
- Unqualified XA switch load files

## API exits:

For a server, use the Exits queue manager properties page from the MQ Explorer, or the ApiExitLocal stanza in the qm.ini file to identify API exit routines for a queue manager. For a client modify the ApiExitLocal stanza in the mqclient.ini file to identify API exit routines for a queue manager.

On Windows systems, you can also use the amqmdain command to change the entries for API exits. (To identify API exit routines for all queue managers, you use the ApiExitCommon and ApiExitTemplate stanzas, as described in “API exits” on page 780.)

Note, that for the API exit to work correctly, the message from the server must be sent to the client unconverted. After the API exit has processed the message, the message must then be converted on the client. This, therefore, requires that you have installed all conversion exits on the client.

For a complete description of the attributes for these stanzas, see Configuring API exits.

## Queue manager error logs

Use the Extended queue manager properties page from the IBM MQ Explorer, or the QMErrorLog stanza in the qm.ini file to tailor the operation and contents of queue manager error logs.


**Attention:** You can use IBM MQ Explorer to make the changes, only if you are using a local queue manager on the Windows platform.

### **ErrorLogSize=** *maxsize*

Specifies the size of the queue manager error log at which it is copied to the backup. *maxsize* must be in the range 32768 through 2147483648 bytes. If **ErrorLogSize** is not specified, the default value of 2097152 bytes (2 MB) is used.

### **ExcludeMessage=** *msgIds*


Specifies messages that are not to be written to the queue manager error log. If your IBM MQ system is heavily used, with many channels stopping and starting, a large number of information messages are sent to the z/OS console and hardcopy log. The IBM MQ - IMS bridge and buffer manager might also produce a large number of information messages, so excluding messages prevents you from receiving a large number of messages if you require it. *msgIds* contain a comma-separated list of message id's from the following:

- 5211 - Maximum property name length exceeded.
- 5973 - Distributed publish/subscribe subscription inhibited
- 5974 - Distributed publish/subscribe publication inhibited
- 6254 - The system could not dynamically load shared library
-  7163 - Job started message ( IBM i only)
- 7234 - Number of messages loaded
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9489 - SVRCONN max instances limit exceeded
- 9490 - SVRCONN max instances per client limit exceeded
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel
- 9545 - Disconnect interval expired
- 9558 - Remote Channel is not available
- 9637 - Channel is lacking a certificate
- 9776 - Channel was blocked by user ID
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address
- 9999 - Channel program ended abnormally

### **SuppressMessage=** *msgIds*

Specifies messages that are written to the queue manager error log once only in a specified time interval. If your IBM MQ system is heavily used, with many channels stopping and starting, a large number of information messages are sent to the z/OS console and hardcopy log. The IBM MQ - IMS bridge and buffer manager might also produce a large number of information messages, so suppressing messages prevents you from receiving a number of repeating messages

if you require it. The time interval is specified by **SuppressInterval** . *msgIds* contain a comma-separated list of message id's from the following:

- 5211 - Maximum property name length exceeded.
- 5973 - Distributed publish/subscribe subscription inhibited
- 5974 - Distributed publish/subscribe publication inhibited
- 6254 - The system could not dynamically load shared library
-  7163 - Job started message ( IBM i only)
- 7234 - Number of messages loaded
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9489 - SVRCONN max instances limit exceeded
- 9490 - SVRCONN max instances per client limit exceeded
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel
- 9545 - Disconnect interval expired
- 9558 - Remote Channel is not available
- 9637 - Channel is lacking a certificate
- 9776 - Channel was blocked by user ID
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address
- 9999 - Channel program ended abnormally

If the same message id is specified in both **SuppressMessage** and **ExcludeMessage** , the message is excluded.

**SuppressInterval=** *length*

Specifies the time interval, in seconds, in which messages specified in **SuppressMessage** are written to the queue manager error log once only. *length* must be in the range 1 through 86400 seconds. If **SuppressInterval** is not specified, the default value of 30 seconds is used.

## Queue manager default bind type

Use the Extended queue manager properties page from the IBM MQ Explorer, or the Connection stanza in the qm.ini file to specify the default bind type. Note that you must create a Connection stanza if you need one.

### DefaultBindType= SHARED | ISOLATED

If DefaultBindType is set to ISOLATED, applications and the queue manager run in separate processes, and no resources are shared between them.

If DefaultBindType is set to SHARED, applications and the queue manager run in separate processes, but some resources are shared between them.

The default is SHARED.

## SSL and TLS stanza of the queue manager configuration file

Use the SSL stanza of the queue manager configuration file to configure SSL or TLS channels on your queue manager.

### Online Certificate Status Protocol (OCSP)

A certificate can contain an AuthorityInfoAccess extension. This extension specifies a server to be contacted through Online Certificate Status Protocol (OCSP). To allow SSL or TLS channels on your queue manager to use AuthorityInfoAccess extensions, ensure that the OCSP server named in them is available, is correctly configured, and is accessible over the network. For more information, see Working with revoked certificates.

### CrIDistributionPoint (CDP)

A certificate can contain a CrIDistributionPoint extension. This extension contains a URL which identifies both the protocol used to download a certificate revocation list (CRL) and also the server to be contacted.

If you want to allow SSL or TLS channels on your queue manager to use CrIDistributionPoint extensions, ensure that the CDP server named in them is available, correctly configured, and accessible over the network.

## The SSL Stanza

Use the SSL stanza in the qm.ini file to configure how SSL or TLS channels on your queue manager attempts to use the following facilities, and how they react if problems occur when using them.

In each of the following cases, if the value supplied is not one of the valid values listed, then the default value is taken. No error messages are written mentioning that an invalid value is specified.

### CDPCheckExtensions= YES | NO

CDPCheckExtensions specifies whether SSL or TLS channels on this queue manager try to check CDP servers that are named in CrIDistributionPoint certificate extensions.

- YES: SSL or TLS channels try to check CDP servers to determine whether a digital certificate is revoked.
- NO: SSL or TLS channels do not try to check CDP servers. This value is the default.

### OCSPAAuthentication= REQUIRED | WARN | OPTIONAL

OCSPAAuthentication specifies the action to be taken when a revocation status cannot be determined from an OCSP server.

If OCSP checking is enabled, an SSL or TLS channel program attempts to contact an OCSP server.

If the channel program is unable to contact any OCSP servers, or if no server can provide the revocation status of the certificate, then the value of the OCSPAAuthentication parameter is used.

- **REQUIRED:** Failure to determine the revocation status causes the connection to be closed with an error. This value is the default.
- **WARN:** Failure to determine the revocation status causes a warning message to be written in the queue manager error log, but the connection is allowed to proceed.
- **OPTIONAL:** Failure to determine the revocation status allows the connection to proceed silently. No warnings or errors are given.

**OCSPCheckExtensions= YES | NO**

OCSPCheckExtensions specifies whether SSL and TLS channels on this queue manager try to check OCSP servers that are named in AuthorityInfoAccess certificate extensions.

- **YES:** SSL and TLS channels try to check OCSP servers to determine whether a digital certificate is revoked. This value is the default.
- **NO:** SSL and TLS channels do not try to check OCSP servers.

**SSLHTTPProxyName= *string***

The string is either the host name or network address of the HTTP Proxy server that is to be used by GSKit for OCSP checks. This address can be followed by an optional port number, enclosed in parentheses. If you do not specify the port number, the default HTTP port, 80, is used. On the HP-UX PA-RISC and Sun Solaris SPARC platforms, and for 32-bit clients on AIX, the network address can only be an IPv4 address; on other platforms it can be an IPv4 or IPv6 address.

This attribute might be necessary if, for example, a firewall prevents access to the URL of the OCSP responder.

## Exit properties

Use the Cluster queue manager properties page from the MQ Explorer, or the ExitPropertiesLocal stanza in the qm.ini file, to specify information about exit properties on a queue manager. Alternatively, you can set it using the **amqmdain** command.

By default, this setting is inherited from the CLWLMode attribute in the ExitProperties stanza of the machine-wide configuration (described in “Exit properties” on page 776 ). Change this setting only if you want to configure this queue manager in a different way. This value can be overridden for individual queue managers using the cluster workload mode attribute on the Cluster queue manager properties page.

**CLWLMode= SAFE | FAST**

The cluster workload (CLWL) exit allows you to specify which cluster queue in the cluster to open in response to an MQI call (for example, **MQOPEN** , **MQPUT** ). The CLWL exit runs either in FAST mode or SAFE mode depending on the value you specify on the CLWLMode attribute. If you omit the CLWLMode attribute, the cluster workload exit runs in SAFE mode.

**SAFE**

Run the CLWL exit in a separate process from the queue manager. This is the default.

If a problem arises with the user-written CLWL exit when running in SAFE mode, the following happens:

- The CLWL server process (amqzlw0) fails.
- The queue manager restarts the CLWL server process.
- The error is reported to you in the error log. If an MQI call is in progress, you receive notification in the form of a return code.

The integrity of the queue manager is preserved.

**Note:** Running the CLWL exit in a separate process can affect performance.

**FAST**

Run the cluster exit inline in the queue manager process.



Specifying this option improves performance by avoiding the process switching costs associated with running in SAFE mode, but does so at the expense of queue manager integrity. You should only run the CLWL exit in FAST mode if you are convinced that there are **no** problems with your CLWL exit, and you are particularly concerned about performance.

If a problem arises when the CLWL exit is running in FAST mode, the queue manager will fail and you run the risk of the integrity of the queue manager being compromised.

## Subpool

This stanza is created by IBM MQ. Do not change it.

The stanza Subpool, and attribute ShortSubpoolName within that stanza, are written automatically by IBM MQ when you create a queue manager. IBM MQ chooses a value for ShortSubpoolName. Do not alter this value.

The name corresponds to a directory and symbolic link created inside the /var/mqm/sockets directory, which IBM MQ uses for internal communications between its running processes.

## Filesystem


From IBM WebSphere MQ Version 7.0 onwards, directory and file systems permissions are more restricted. By default, only members of the mqm group can write directly to error log files and First Failure Data Capture files. You can use the **Filesystem** stanza to allow users who are not members of the mqm group to access error directories and files.

To allow users who are not members of the mqm group, on IBM i, members of the QMQMADM group, to access error directories and files, you must set:

**ValidateAuth=**

No

Note that the text is case sensitive.

 On IBM i, you must also set the authority for the additional users to \*PUBLIC.

**Note:** IBM MQ does not support the addition of users to error logs.

You can use this to extend access, by changing the group ownership of the directory and using setgid permissions. For example, to widen access to include members of a group called mqerrors, use the following:

```
chgrp mqerrors /var/mqm/errors
chgrp mqerrors /var/mqm/qmgrs/<QMNAME>/errors
chmod 6770 /var/mqm/qmgrs/<QMNAME>/errors
```

This causes all files within these directories to be created with mqerrors ownership, rather than mqm ownership. Hence, extending access to the members of the mqerrors group.

This approach does not provide o+r permissions on the actual files. Alternatively, a cron job (running under mqm) could periodically change the permissions of the files within these directories, to provide o+r permissions





---

## Configuring distributed queuing

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

Before reading this section it is helpful to have an understanding of channels, queues, and the other concepts introduced in Distributed queuing and clusters.

Use the information in the following links to connect your applications using distributed queuing:

- “How to send a message to another queue manager” on page 826
- “Triggering channels” on page 848
- “Safety of messages” on page 846
- “IBM MQ distributed queuing techniques”
- “Introduction to distributed queue management” on page 823
-    “Monitoring and controlling channels on Windows, UNIX and Linux platforms” on page 856
-  “Monitoring and controlling channels on IBM i” on page 879

### Related concepts:

 “Customizing IBM MQ for z/OS” on page 1198


Use this topic as a step by step guide for customizing your IBM MQ system.

“Configuring connections between the server and client” on page 696


To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

“Changing IBM MQ and queue manager configuration information” on page 755

Change the behavior of IBM MQ or an individual queue manager to suit the needs of your installation.

 “Configuring queue managers on z/OS” on page 1194

Use these instructions to configure queue managers on IBM MQ for z/OS.

 “Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

### Related tasks:




“Configuring a queue manager cluster” on page 902

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

## IBM MQ distributed queuing techniques

The subtopics in this section describe techniques that are of use when planning channels. These subtopics describe techniques to help you plan how to connect your queue managers together, and manage the flow of messages between your applications.

For message channel planning examples, see:

- Message channel planning example for distributed platforms
-  Message channel planning example for IBM MQ for IBM i
-  Message channel planning example for z/OS
-  Message channel planning example for z/OS using queue-sharing groups

**Related concepts:**

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

**Related information:**

Channels


Introduction to message queuing

Distributed queuing and clusters

Example configuration information

**Message flow control**

Message flow control is a task that involves the setting up and maintenance of message routes between queue managers. It is important for routes that multi-hop through many queue managers. This section describes how you use queues, alias queue definitions, and message channels on your system to achieve message flow control.

You control message flow using a number of techniques that were introduced in “Configuring distributed queuing” on page 802. If your queue manager is in a cluster, message flow is controlled using different techniques, as described in “Message flow control.”  If your queue managers are in a queue sharing group and intra-group queuing (IGQ) is enabled, then the message flow can be controlled by IGQ agents. These agents are described in Intra-group queuing.

You can use the following objects to achieve message flow control:

- Transmission queues
- Message channels
- Remote queue definition
- Queue manager alias definition
- Reply-to queue alias definition

The queue manager and queue objects are described in Object types. Message channels are described in Distributed queuing components. The following techniques use these objects to create message flows in your system:

- Putting messages to remote queues
- Routing by way of particular transmission queues
- Receiving messages
- Passing messages through your system
- Separating message flows
- Switching a message flow to another destination
- Resolving the reply-to queue name to an alias name

**Note**

All the concepts described in this section are relevant for all nodes in a network, and include sending and receiving ends of message channels. For this reason, only one node is illustrated in most examples. The exception is where the example requires explicit cooperation by the administrator at the other end of a message channel.

Before proceeding to the individual techniques, it is useful to recap on the concepts of name resolution and the three ways of using remote queue definitions. See Distributed queuing and clusters.

**Related concepts:**

“Queue names in transmission header”

Destination queue names travel with the message in the transmission header until the destination queue has been reached.

“How to create queue manager and reply-to aliases”

This topic explains the three ways that you can create a remote queue definition.

**Queue names in transmission header:**

Destination queue names travel with the message in the transmission header until the destination queue has been reached.

The queue name used by the application, the logical queue name, is resolved by the queue manager to the destination queue name. In other words, the physical queue name. This destination queue name travels with the message in a separate data area, the transmission header, until the destination queue has been reached. The transmission header is then stripped off.

You change the queue manager part of this queue name when you create parallel classes of service. Remember to return the queue manager name to the original name when the end of the class-of-service diversion has been reached.

**How to create queue manager and reply-to aliases:**

This topic explains the three ways that you can create a remote queue definition.

The remote queue definition object is used in three different ways. Table 124 on page 805 explains how to define each of the three ways:

- Using a remote queue definition to redefine a local queue name.

The application provides only the queue name when opening a queue, and this queue name is the name of the remote queue definition.

The remote queue definition contains the names of the target queue and queue manager. Optionally, the definition can contain the name of the transmission queue to be used. If no transmission queue name is provided, the queue manager uses the queue manager name, taken from the remote queue definition, for the transmission queue name. If a transmission queue of this name is not defined, but a default transmission queue is defined, the default transmission queue is used.

- Using a remote queue definition to redefine a queue manager name.

The application, or channel program, provides a queue name together with the remote queue manager name when opening the queue.

If you have provided a remote queue definition with the same name as the queue manager name, and you have left the queue name in the definition blank, then the queue manager substitutes the queue manager name in the open call with the queue manager name in the definition.

In addition, the definition can contain the name of the transmission queue to be used. If no transmission queue name is provided, the queue manager takes the queue manager name, taken from the remote queue definition, for the transmission queue name. If a transmission queue of this name is not defined, but a default transmission queue is defined, the default transmission queue is used.

- Using a remote queue definition to redefine a reply-to queue name.

Each time an application puts a message to a queue, it can provide the name of a reply-to queue for answer messages but with the queue manager name blank.

If you provide a remote queue definition with the same name as the reply-to queue then the local queue manager replaces the reply-to queue name with the queue name from your definition.

You can provide a queue manager name in the definition, but not a transmission queue name.

Table 124. Three ways of using the remote queue definition object

| Usage                                     | Queue manager name            | Queue name   | Transmission queue name |
|-------------------------------------------|-------------------------------|--------------|-------------------------|
| 1. Remote queue definition (on OPEN call) |                               |              |                         |
| Supplied in the call                      | blank or local QM             | (*) required | not applicable          |
| Supplied in the definition                | required                      | required     | optional                |
| 2. Queue manager alias (on OPEN call)     |                               |              |                         |
| Supplied in the call                      | (*) required and not local QM | required     | not applicable          |
| Supplied in the definition                | required                      | blank        | optional                |
| 3. Reply-to queue alias (on PUT call)     |                               |              |                         |
| Supplied in the call                      | blank                         | (*) required | not applicable          |
| Supplied in the definition                | optional                      | optional     | blank                   |

**Note:** (\*) means that this name is the name of the definition object

For a formal description, see Queue name resolution.

### Putting messages on remote queues

You can use remote queue definition objects to resolve a queue name to a transmission queue to an adjacent queue manager.

In a distributed-queuing environment, a transmission queue and channel are the focal point for all messages to a location whether the messages originate from applications in your local system, or arrive through channels from an adjacent system. Figure 86 on page 806 shows an application placing messages on a logical queue named 'QA\_norm'. The name resolution uses the remote queue definition 'QA\_norm' to select the transmission queue QMB. It then adds a transmission header to the messages stating 'QA\_norm at QMB'.

Messages arriving from the adjacent system on 'Channel\_back' have a transmission header with the physical queue name 'QA\_norm at QMB', for example. These messages are placed unchanged on transmission queue QMB.

The channel moves the messages to an adjacent queue manager.

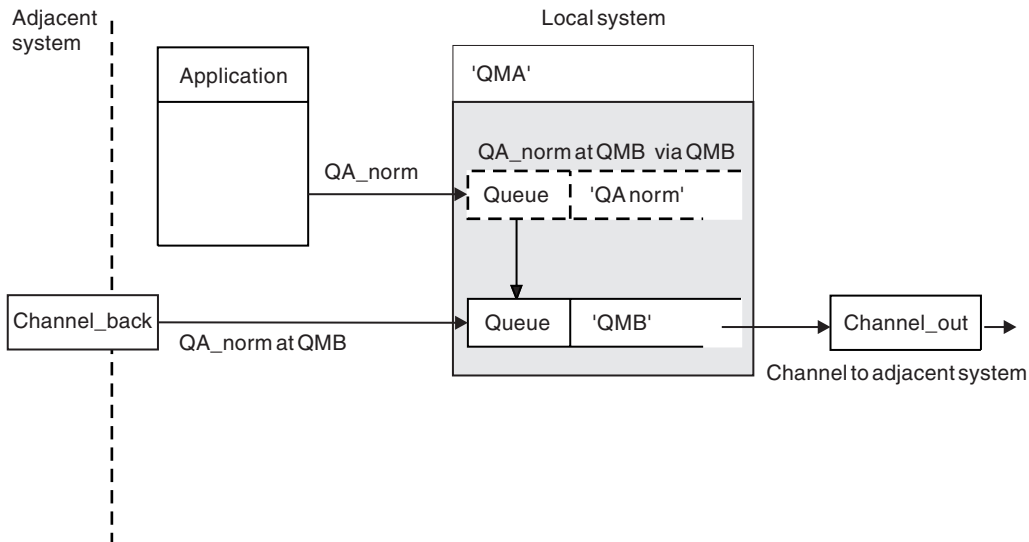


Figure 86. A remote queue definition is used to resolve a queue name to a transmission queue to an adjacent queue manager. Note: The dashed outline represents a remote queue definition. This queue is not a real queue, but a name alias that is controlled as though it were a real queue.

If you are the IBM MQ system administrator, you must:

- Define the message channel from the adjacent system
- Define the message channel to the adjacent system
- Create the transmission queue QMB
- Define the remote queue object 'QA\_norm' to resolve the queue name used by applications to the destination queue name, destination queue manager name, and transmission queue name

In a clustering environment, you only need to define a cluster-receiver channel at the local queue manager. You do not need to define a transmission queue or a remote queue object. See Clusters.

### More about name resolution

The effect of the remote queue definition is to define a physical destination queue name and queue manager name. These names are put in the transmission headers of messages.

Incoming messages from an adjacent system have already had this type of name resolution carried out by the original queue manager. Therefore they have the transmission header showing the physical destination queue name and queue manager name. These messages are unaffected by remote queue definitions.

## Choosing the transmission queue

You can use a remote queue definition to allow a different transmission queue to send messages to the same adjacent queue manager.

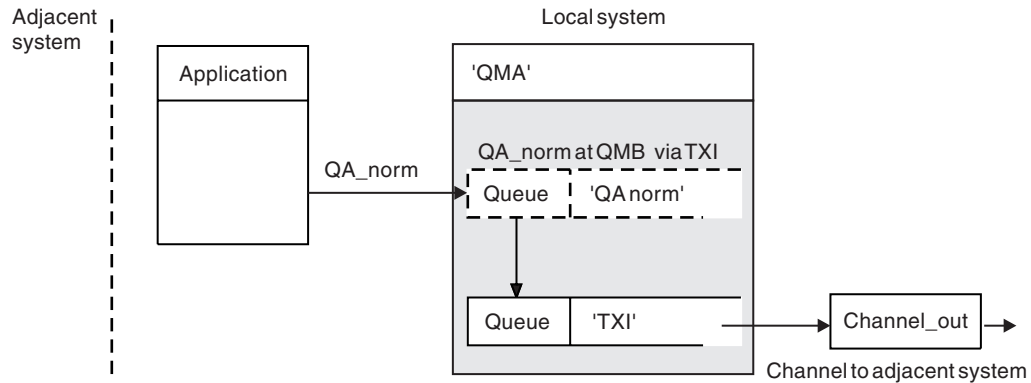


Figure 87. The remote queue definition allows a different transmission queue to be used

In a distributed-queuing environment, when you need to change a message flow from one channel to another, use the same system configuration as shown in Figure 86 on page 806 in “Putting messages on remote queues” on page 805. Figure 87 in this topic shows how you use the remote queue definition to send messages over a different transmission queue, and therefore over a different channel, to the same adjacent queue manager.

For the configuration shown in Figure 87, you must provide the remote queue object 'QA\_norm', and the transmission queue 'TX1'. You must provide 'QA\_norm' to choose the 'QA\_norm' queue at the remote queue manager, the transmission queue 'TX1', and the queue manager 'QMB\_priority'. Specify 'TX1' in the definition of the channel adjacent to the system.

Messages are placed on transmission queue 'TX1' with a transmission header containing 'QA\_norm at QMB\_priority', and are sent over the channel to the adjacent system.

The channel\_back has been left out of this illustration because it would need a queue manager alias.

In a clustering environment, you do not need to define a transmission queue or a remote queue definition. For more information, see “Defining cluster queues” on page 903.

## Receiving messages

You can configure the queue manager to receive messages from other queue managers. You must ensure that unintentional name resolution does not occur.

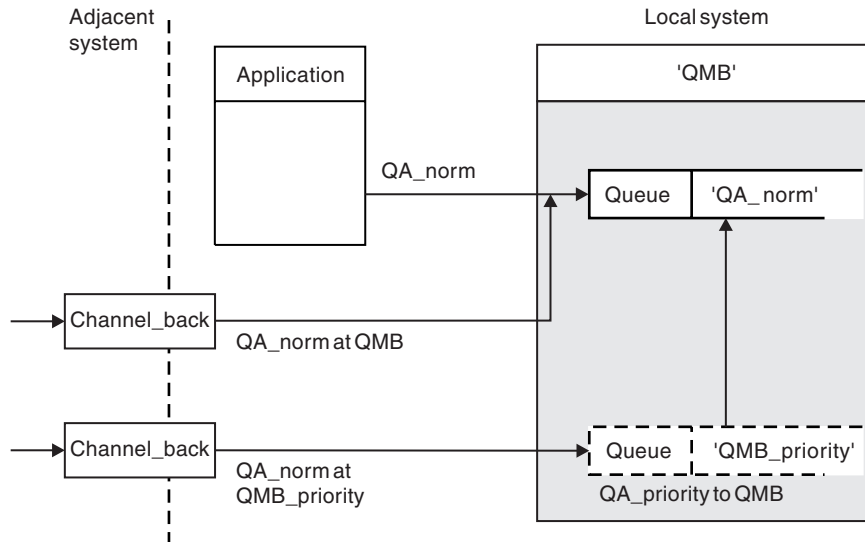


Figure 88. Receiving messages directly, and resolving alias queue manager name

As well as arranging for messages to be sent, the system administrator must also arrange for messages to be received from adjacent queue managers. Received messages contain the physical name of the destination queue manager and queue in the transmission header. They are treated the same as messages from a local application that specifies both queue manager name and queue name. Because of this treatment, you need to ensure that messages entering your system do not have an unintentional name resolution carried out. See Figure 88 for this scenario.

For this configuration, you must prepare:

- Message channels to receive messages from adjacent queue managers
- A queue manager alias definition to resolve an incoming message flow, 'QMB\_priority', to the local queue manager name, 'QMB'
- The local queue, 'QA\_norm', if it does not exist

## Receiving alias queue manager names

The use of the queue manager alias definition in this illustration has not selected a different destination queue manager. Messages passing through this local queue manager and addressed to 'QMB\_priority' are intended for queue manager 'QMB'. The alias queue manager name is used to create the separate message flow.



## Passing messages through your system

You can pass messages through your system in three ways - using the location name, using an alias for the queue manager, or selecting a transmission queue.

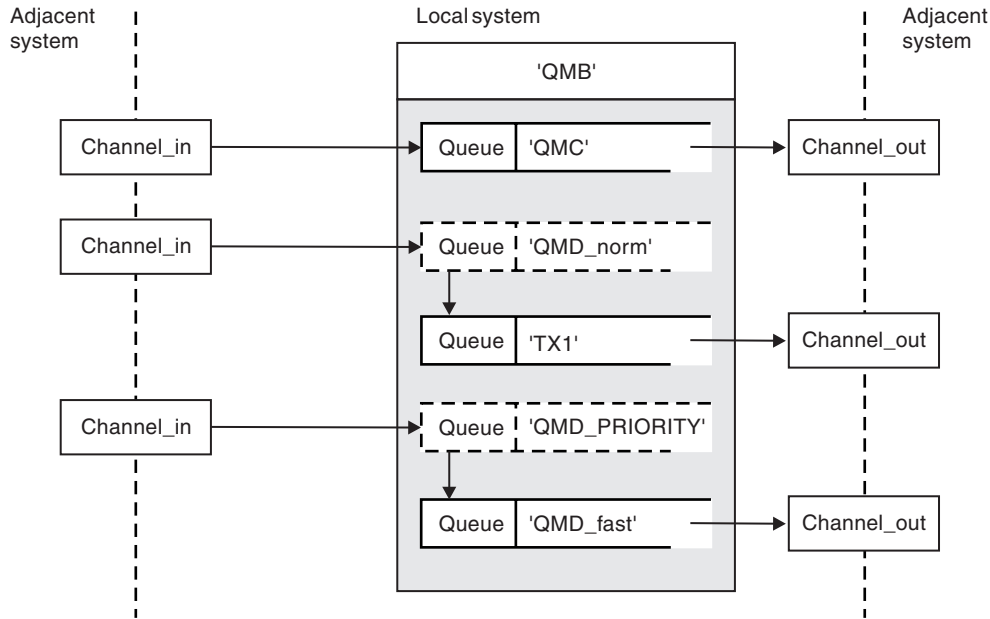


Figure 89. Three methods of passing messages through your system

The technique shown in Figure 88 on page 808 in “Receiving messages” on page 808, showed how an alias flow is captured. Figure 89 illustrates the ways networks are built up by bringing together the techniques previously described.

The configuration shows a channel delivering three messages with different destinations:

1. QB at QMC
2. QB at QMD\_norm
3. QB at QMD\_PRIORITY

You must pass the first message flow through your system unchanged. You must pass the second message flow through a different transmission queue and channel. For the second message flow you must also resolve messages for the alias queue manager name QMD\_norm to the queue manager QMD. The third message flow chooses a different transmission queue without any other change.

In a clustering environment, messages are passed through a cluster transmission queue. Normally a single transmission queue, SYSTEM.CLUSTER.TRANSMIT.QUEUE, transfers all messages to all queue managers in all clusters that the queue manager is a member of; see A cluster of queue managers. You can define separate transmission queues for all or some of the queue managers in the clusters that the queue manager is a member of.

The following methods describe techniques applicable to a distributed-queuing environment.

### Use these methods

For these configurations, you must prepare the:

- Input channel definitions

- Output channel definitions
- Transmission queues:
  - QMC
  - TX1
  - QMD\_fast
- Queue manager alias definitions:
  - QMD\_norm with QMD\_norm to QMD through TX1
  - QMD\_PRIORITY with QMD\_PRIORITY to QMD\_PRIORITY through QMD\_fast

**Note:** None of the message flows shown in the example changes the destination queue. The queue manager name aliases provide separation of message flows.

### **Method 1: Use the incoming location name**

You are going to receive messages with a transmission header containing another location name, such as QMC. The simplest configuration is to create a transmission queue with that name, QMC. The channel that services the transmission queue delivers the message unchanged to the next destination.

### **Method 2: Use an alias for the queue manager**

The second method is to use the queue manager alias object definition, but specify a new location name, QMD, and a particular transmission queue, TX1. This action:

- Terminates the alias message flow setup by the queue manager name alias QMD\_norm, that is, the named class of service QMD\_norm.
- Changes the transmission headers on these messages from QMD\_norm to QMD.

### **Method 3: Select a transmission queue**

The third method is to have a queue manager alias object defined with the same name as the destination location, QMD\_PRIORITY. Use the queue manager alias definition to select a particular transmission queue, QMD\_fast, and therefore another channel. The transmission headers on these messages remain unchanged.

### **Separating message flows**

You can use a queue manager alias to create separate message flows to send messages to the same queue manager.

In a distributed-queuing environment, the need to separate messages to the same queue manager into different message flows can arise for a number of reasons. For example:

- You might need to provide a separate flow for large, medium, and small messages. This need also applies in a clustering environment and, in this case, you can create clusters that overlap. There are a number of reasons you might do so, for example:
  - To allow different organizations to have their own administration.
  - To allow independent applications to be administered separately.
  - To create a class of service. For example, you could have a cluster called STAFF that is a subset of the cluster called STUDENTS. When you put a message to a queue advertised in the STAFF cluster, a restricted channel is used. When you put a message to a queue advertised in the STUDENTS cluster, either a general channel or a restricted channel can be used.
  - To create test and production environments.
- It might be necessary to route incoming messages by different paths from the path of the locally generated messages.
- Your installation might require to schedule the movement of messages at certain times (for example, overnight) and the messages then need to be stored in reserved queues until scheduled.

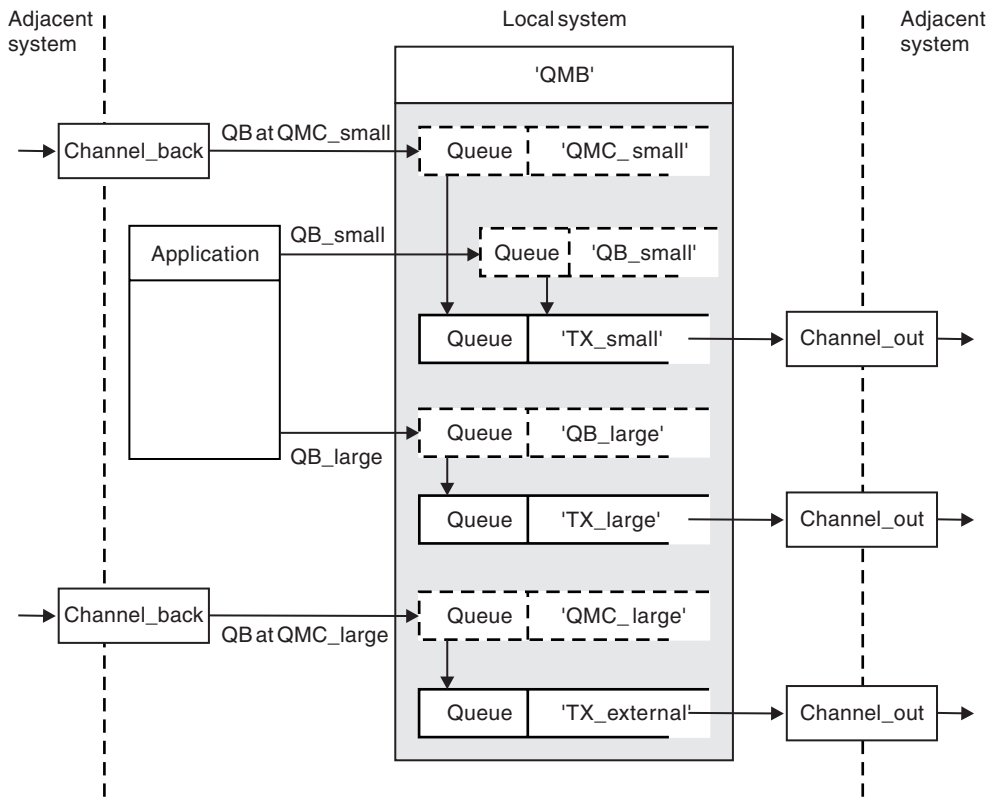


Figure 90. Separating messages flows

In the example shown in Figure 90, the two incoming flows are to alias queue manager names 'QMC\_small' and 'QMC\_large'. You provide these flows with a queue manager alias definition to capture these flows for the local queue manager. You have an application addressing two remote queues and you need these message flows to be kept separate. You provide two remote queue definitions that specify the same location, 'QMC', but specify different transmission queues. This definition keeps the flows separate, and nothing extra is needed at the far end as they have the same destination queue manager name in the transmission headers. You provide:

- The incoming channel definitions
- The two remote queue definitions QB\_small and QB\_large
- The two queue manager alias definitions QMC\_small and QMC\_large
- The three sending channel definitions
- Three transmission queues: TX\_small, TX\_large, and TX\_external

### Coordination with adjacent systems

When you use a queue manager alias to create a separate message flow, you need to coordinate this activity with the system administrator at the remote end of the message channel to ensure that the corresponding queue manager alias is available there.

## Concentrating messages to diverse locations

You can concentrate messages destined for various locations on to a single channel.

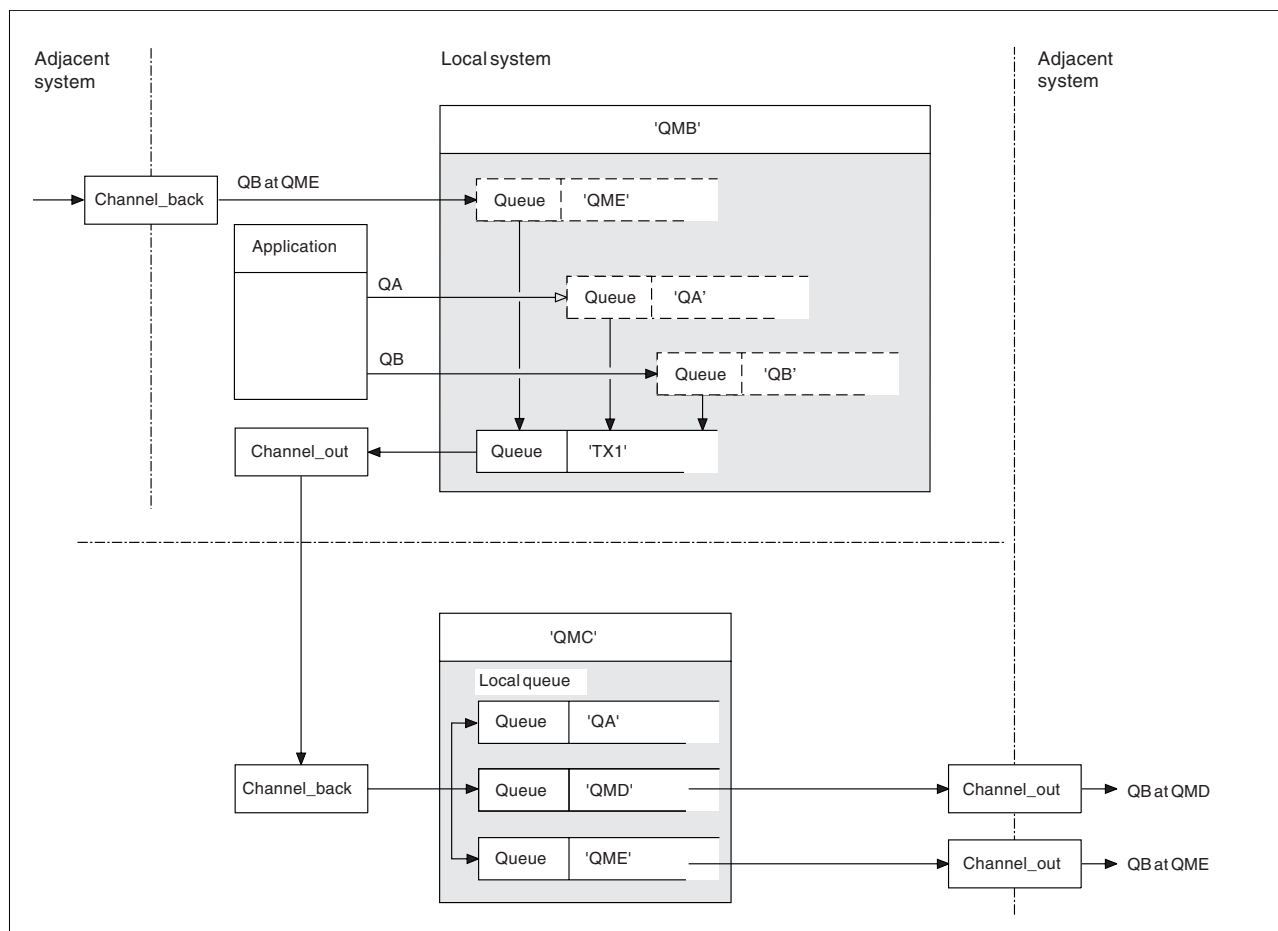


Figure 91. Combining message flows on to a channel

Figure 91 illustrates a distributed-queuing technique for concentrating messages that are destined for various locations on to one channel. Two possible uses would be:

- Concentrating message traffic through a gateway
- Using wide bandwidth highways between nodes

In this example, messages from different sources, local and adjacent, and having different destination queues and queue managers, are flowed through transmission queue 'TX1' to queue manager QMC. Queue manager QMC delivers the messages according to the destinations. One set to a transmission queue 'QMD' for onward transmission to queue manager QMD. Another set to a transmission queue 'QME' for onward transmission to queue manager QME. Other messages are put on the local queue 'QA'.

You must provide:

- Channel definitions
- Transmission queue TX1
- Remote queue definitions:
  - QA with 'QA at QMC through TX1'
  - QB with 'QB at QMD through TX1'
- Queue manager alias definition:

- QME with 'QME through TX1'

The complementary administrator who is configuring QMC must provide:

- Receiving channel definition with the same channel name
- Transmission queue QMD with associated sending channel definition
- Transmission queue QME with associated sending channel definition
- Local queue object QA.

## Diverting message flows to another destination

You can redefine the destination of certain messages using queue manager aliases and transmission queues.

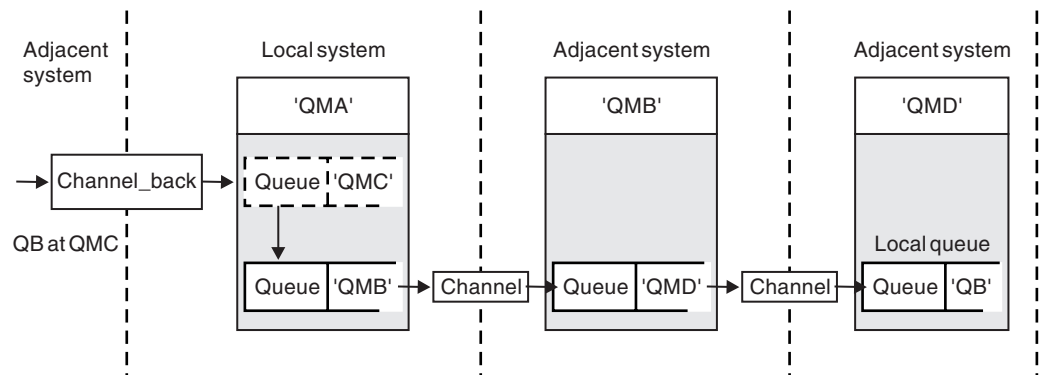


Figure 92. Diverting message streams to another destination

Figure 92 illustrates how you can redefine the destination of certain messages. Incoming messages to QMA are destined for 'QB at QMC'. They normally arrive at QMA and be placed on a transmission queue called QMC which has been part of a channel to QMC. QMA must divert the messages to QMD, but is able to reach QMD only over QMB. This method is useful when you need to move a service from one location to another, and allow subscribers to continue to send messages on a temporary basis until they have adjusted to the new address.

The method of rerouting incoming messages destined for a certain queue manager to a different queue manager uses:

- A queue manager alias to change the destination queue manager to another queue manager, and to select a transmission queue to the adjacent system
- A transmission queue to serve the adjacent queue manager
- A transmission queue at the adjacent queue manager for onward routing to the destination queue manager

You must provide:

- Channel\_back definition
- Queue manager alias object definition QMC with QB at QMD through QMB
- Channel\_out definition
- The associated transmission queue QMB

The complementary administrator who is configuring QMB must provide:

- The corresponding channel\_back definition
- The transmission queue, QMD

- The associated channel definition to QMD

You can use aliases within a clustering environment. For information, see “Queue-manager aliases and clusters” on page 984.

## Sending messages to a distribution list

You can use a single MQPUT call to have an application send a message to several destinations.

In IBM MQ on all platforms except z/OS, an application can send a message to several destinations with a single MQPUT call. You can do so in both a distributed-queuing environment and a clustering environment. You have to define the destinations in a distribution list, as described in Distribution lists.

Not all queue managers support distribution lists. When an MCA establishes a connection with a partner, it determines whether the partner supports distribution lists and sets a flag on the transmission queue accordingly. If an application tries to send a message that is destined for a distribution list but the partner does not support distribution lists, the sending MCA intercepts the message and puts it onto the transmission queue once for each intended destination.

A receiving MCA ensures that messages sent to a distribution list are safely received at all the intended destinations. If any destinations fail, the MCA establishes which ones have failed. It then can generate exception reports for them and can try to send the messages to them again.

## Reply-to queue

You can create a complete remote queue processing loop using a reply-to queue.

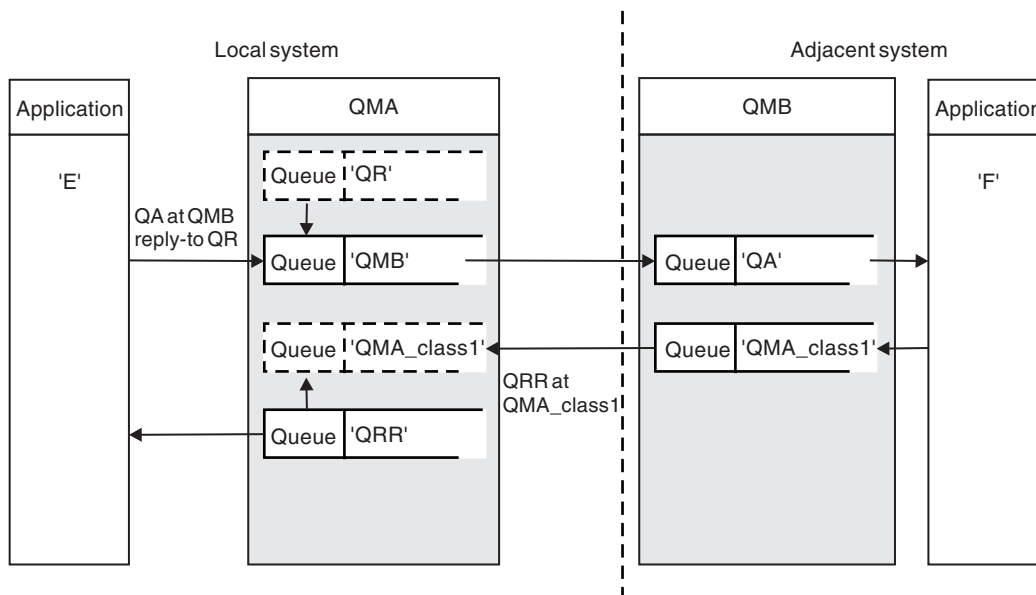


Figure 93. Reply-to queue name substitution during PUT call

A complete remote queue processing loop using a reply-to queue is shown in Figure 93. This loop applies in both a distributed-queuing environment and a clustering environment. The details are as shown in Table 128 on page 822.

The application opens QA at QMB and puts messages on that queue. The messages are given a reply-to queue name of QR, without the queue manager name being specified. Queue manager QMA finds the reply-to queue object QR and extracts from it the alias name of QRR and the queue manager name QMA\_class1. These names are put into the reply-to fields of the messages.

Reply messages from applications at QMB are addressed to QRR at QMA\_class1. The queue manager alias name definition QMA\_class1 is used by the queue manager to flow the messages to itself, and to queue QRR.

This scenario depicts the way you give applications the facility to choose a class of service for reply messages. The class is implemented by the transmission queue QMA\_class1 at QMB, together with the queue manager alias definition, QMA\_class1 at QMA. In this way, you can change an application's reply-to queue so that the flows are segregated without involving the application. The application always chooses QR for this particular class of service. You have the opportunity to change the class of service with the reply-to queue definition QR.

You must create:

- Reply-to queue definition QR
- Transmission queue object QMB
- Channel\_out definition
- Channel\_back definition
- Queue manager alias definition QMA\_class1
- Local queue object QRR, if it does not exist

The complementary administrator at the adjacent system must create:

- Receiving channel definition
- Transmission queue object QMA\_class1
- Associated sending channel
- Local queue object QA.

Your application programs use:

- Reply-to queue name QR in put calls
- Queue name QRR in get calls

In this way, you can change the class of service as necessary, without involving the application. You change the reply-to alias 'QR', together with the transmission queue 'QMA\_class1' and queue manager alias 'QMA\_class1'.

If no reply-to alias object is found when the message is put on the queue, the local queue manager name is inserted in the blank reply-to queue manager name field. The reply-to queue name remains unchanged.

### **Name resolution restriction**

Because the name resolution has been carried out for the reply-to queue at 'QMA' when the original message was put, no further name resolution is allowed at 'QMB'. The message is put with the physical name of the reply-to queue by the replying application.

The applications must be aware that the name they use for the reply-to queue is different from the name of the actual queue where the return messages are to be found.

For example, when two classes of service are provided for the use of applications with reply-to queue alias names of 'C1\_alias', and 'C2\_alias', the applications use these names as reply-to queue names in the message put calls. However, the applications actually expect messages to appear in queues 'C1' for 'C1\_alias' and 'C2' for 'C2\_alias'.

However, an application is able to make an inquiry call on the reply-to alias queue to check for itself the name of the real queue it must use to get the reply messages.

### Related concepts:

“How to create queue manager and reply-to aliases” on page 804

This topic explains the three ways that you can create a remote queue definition.

“Reply-to queue alias example”

This example illustrates the use of a reply-to alias to select a different route (transmission queue) for returned messages. The use of this facility requires the reply-to queue name to be changed in cooperation with the applications.

“How the example works” on page 818

An explanation of the example and how the queue manager uses the reply-to queue alias.

“Reply-to queue alias walk-through” on page 818

A walk-through of the process from an application putting a message on a remote queue through to the same application removing the reply message from the alias reply-to queue.

### Reply-to queue alias example:

This example illustrates the use of a reply-to alias to select a different route (transmission queue) for returned messages. The use of this facility requires the reply-to queue name to be changed in cooperation with the applications.

As shown in Figure 94, the return route must be available for the reply messages, including the transmission queue, channel, and queue manager alias.

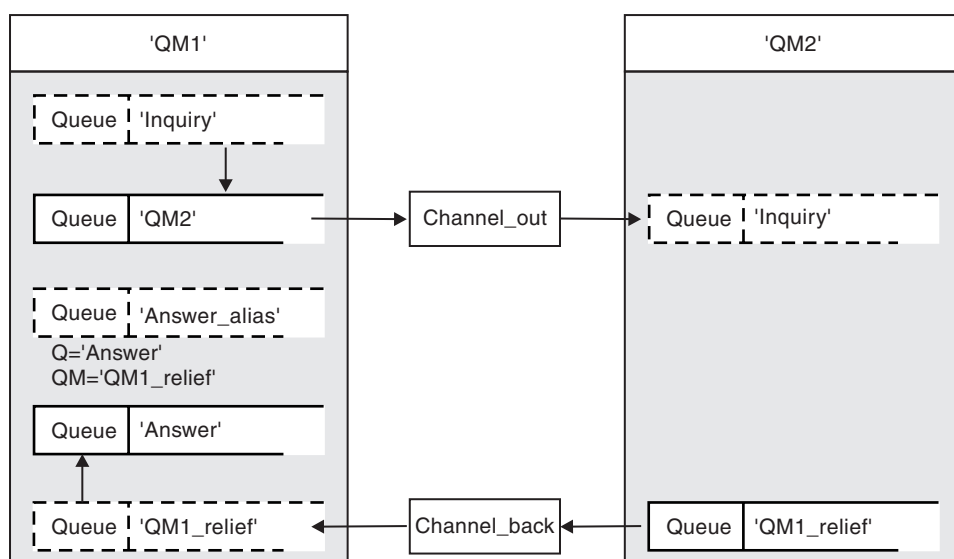


Figure 94. Reply-to queue alias example

This example is for requester applications at 'QM1' that send messages to server applications at 'QM2'. The messages on the server are to be returned through an alternative channel using transmission queue 'QM1\_relief' (the default return channel would be served with a transmission queue 'QM1').

The reply-to queue alias is a particular use of the remote queue definition named 'Answer\_alias'. Applications at QM1 include this name, 'Answer\_alias', in the reply-to field of all messages that they put on queue 'Inquiry'.

Reply-to queue definition 'Answer\_alias' is defined as 'Answer at QM1\_relief'. Applications at QM1 expect their replies to appear in the local queue named 'Answer'.



Server applications at QM2 use the reply-to field of received messages to obtain the queue and queue manager names for the reply messages to the requester at QM1.

### Definitions used in this example at QM1

The IBM MQ system administrator at QM1 must ensure that the reply-to queue 'Answer' is created along with the other objects. The name of the queue manager alias, marked with a '\*', must agree with the queue manager name in the reply-to queue alias definition, also marked with an '\*'.

| Object                   | Definition                |               |
|--------------------------|---------------------------|---------------|
| Local transmission queue | QM2                       |               |
| Remote queue definition  | Object name               | Inquiry       |
|                          | Remote queue manager name | QM2           |
|                          | Remote queue name         | Inquiry       |
|                          | Transmission queue name   | QM2 (DEFAULT) |
| Queue manager alias      | Object name               | QM1_relief *  |
|                          | Queue manager name        | QM1           |
|                          | Queue name                | (blank)       |
| Reply-to queue alias     | Object name               | Answer_alias  |
|                          | Remote queue manager name | QM1_relief *  |
|                          | Remote queue name         | Answer        |

### Put definition at QM1

Applications fill the reply-to fields with the reply-to queue alias name, and leave the queue manager name field blank.

| Field                  | Content      |
|------------------------|--------------|
| Queue name             | Inquiry      |
| Queue manager name     | (blank)      |
| Reply-to queue name    | Answer_alias |
| Reply-to queue manager | (blank)      |

### Definitions used in this example at QM2

The IBM MQ system administrator at QM2 must ensure that the local queue exists for the incoming messages, and that the correctly named transmission queue is available for the reply messages.

| Object             | Definition |
|--------------------|------------|
| Local queue        | Inquiry    |
| Transmission queue | QM1_relief |

### Put definition at QM2

Applications at QM2 retrieve the reply-to queue name and queue manager name from the original message and use them when putting the reply message on the reply-to queue.

| Field              | Content    |
|--------------------|------------|
| Queue name         | Answer     |
| Queue manager name | QM1_relief |

### How the example works:

An explanation of the example and how the queue manager uses the reply-to queue alias.

In this example, requester applications at QM1 always use 'Answer\_alias' as the reply-to queue in the relevant field of the put call. They always retrieve their messages from the queue named 'Answer'.

The reply-to queue alias definitions are available for use by the QM1 system administrator to change the name of the reply-to queue 'Answer', and of the return route 'QM1\_relief'.

Changing the queue name 'Answer' is normally not useful because the QM1 applications are expecting their answers in this queue. However, the QM1 system administrator is able to change the return route (class of service), as necessary.

### How the queue manager uses the reply-to queue alias

Queue manager QM1 retrieves the definitions from the reply-to queue alias when the reply-to queue name, included in the put call by the application, is the same as the reply-to queue alias, and the queue manager part is blank.

The queue manager replaces the reply-to queue name in the put call with the queue name from the definition. It replaces the blank queue manager name in the put call with the queue manager name from the definition.

These names are carried with the message in the message descriptor.

*Table 125. Reply-to queue alias*

| Field name                  | Put call     | Transmission header |
|-----------------------------|--------------|---------------------|
| Reply-to queue name         | Answer_alias | Answer              |
| Reply-to queue manager name | (blank)      | QM1_relief          |

### Reply-to queue alias walk-through:

A walk-through of the process from an application putting a message on a remote queue through to the same application removing the reply message from the alias reply-to queue.

To complete this example, let us look at the process.

1. The application opens a queue named 'Inquiry', and puts messages to it. The application sets the reply-to fields of the message descriptor to:

|                             |                     |
|-----------------------------|---------------------|
| <b>Reply-to queue name</b>  | <b>Answer_alias</b> |
| Reply-to queue manager name | (blank)             |

- Queue manager 'QM1' responds to the blank queue manager name by checking for a remote queue definition with the name 'Answer\_alias'. If none is found, the queue manager places its own name, 'QM1', in the reply-to queue manager field of the message descriptor.
- If the queue manager finds a remote queue definition with the name 'Answer\_alias', it extracts the queue name and queue manager names from the definition (queue name='Answer' and queue manager name='QM1\_relief'). It then puts them into the reply-to fields of the message descriptor.
- The queue manager 'QM1' uses the remote queue definition 'Inquiry' to determine that the intended destination queue is at queue manager 'QM2', and the message is placed on the transmission queue 'QM2'. 'QM2' is the default transmission queue name for messages destined for queues at queue manager 'QM2'.
- When queue manager 'QM1' puts the message on the transmission queue, it adds a transmission header to the message. This header contains the name of the destination queue, 'Inquiry', and the destination queue manager, 'QM2'.
- The message arrives at queue manager 'QM2', and is placed on the 'Inquiry' local queue.
- An application gets the message from this queue and processes the message. The application prepares a reply message, and puts this reply message on the reply-to queue name from the message descriptor of the original message:

|                             |               |
|-----------------------------|---------------|
| <b>Reply-to queue name</b>  | <b>Answer</b> |
| Reply-to queue manager name | QM1_relief    |

- Queue manager 'QM2' carries out the put command. Finding that the queue manager name, 'QM1\_relief', is a remote queue manager, it places the message on the transmission queue with the same name, 'QM1\_relief'. The message is given a transmission header containing the name of the destination queue, 'Answer', and the destination queue manager, 'QM1\_relief'.
- The message is transferred to queue manager 'QM1'. The queue manager, recognizes that the queue manager name 'QM1\_relief' is an alias, extracts from the alias definition 'QM1\_relief' the physical queue manager name 'QM1'.
- Queue manager 'QM1' then puts the message on the queue name contained in the transmission header, 'Answer'.
- The application extracts its reply message from the queue 'Answer'.

## Networking considerations

In a distributed-queuing environment, because message destinations are addressed with just a queue name and a queue manager name, certain rules apply.

- Where the queue manager name is given, and the name is different from the local queue manager name:
  - A transmission queue must be available with the same name. This transmission queue must be part of a message channel moving messages to another queue manager, or
  - A queue manager alias definition must exist to resolve the queue manager name to the same, or another queue manager name, and optional transmission queue, or
  - If the transmission queue name cannot be resolved, and a default transmission queue has been defined, the default transmission queue is used.
- Where only the queue name is supplied, a queue of any type but with the same name must be available on the local queue manager. This queue can be a remote queue definition which resolves to: a transmission queue to an adjacent queue manager, a queue manager name, and an optional transmission queue.

To see how this works in a clustering environment, see Clusters.

► **z/OS** If the queue managers are running in a queue-sharing group (QSG) and intra-group queuing (IGQ) is enabled, you can use the `SYSTEM.QSG.TRANSMIT.QUEUE`. For more information, see Intra-group queuing.

Consider the scenario of a message channel moving messages from one queue manager to another in a distributed-queuing environment.

The messages being moved have originated from any other queue manager in the network, and some messages might arrive that have an unknown queue manager name as destination. This issue can occur when a queue manager name has changed or has been removed from the system, for example.

The channel program recognizes this situation when it cannot find a transmission queue for these messages, and places the messages on your undelivered-message (dead-letter) queue. It is your responsibility to look for these messages and arrange for them to be forwarded to the correct destination. Alternatively, return them to the originator, where the originator can be ascertained.

Exception reports are generated in these circumstances, if report messages were requested in the original message.

### **Name resolution convention**

Name resolution that changes the identity of the destination queue (that is, logical to physical name changing), only occurs once, and only at the originating queue manager.

Subsequent use of the various alias possibilities must only be used when separating and combining message flows.

### **Return routing**

Messages can contain a return address in the form of the name of a queue and queue manager. This return address form can be used in both a distributed-queuing environment and a clustering environment.

This address is normally specified by the application that creates the message. It can be modified by any application that then handles the message, including user exit applications.

Irrespective of the source of this address, any application handling the message might choose to use this address for returning answer, status, or report messages to the originating application.

The way these response messages are routed is not different from the way the original message is routed. You need to be aware that the message flows you create to other queue managers need corresponding return flows.

### **Physical name conflicts**

The destination reply-to queue name has been resolved to a physical queue name at the original queue manager. It must not be resolved again at the responding queue manager.

It is a likely possibility for name conflict problems that can only be prevented by a network-wide agreement on physical and logical queue names.

## Managing queue name translations

When you create a queue manager alias definition or a remote queue definition, the name resolution is carried out for every message carrying that name. This situation must be managed.

This description is provided for application designers and channel planners concerned with an individual system that has message channels to adjacent systems. It takes a local view of channel planning and control.

When you create a queue manager alias definition or a remote queue definition, the name resolution is carried out for every message carrying that name, regardless of the source of the message. To oversee this situation, which might involve large numbers of queues in a queue manager network, you keep tables of:

- The names of source queues and of source queue managers with respect to resolved queue names, resolved queue manager names, and resolved transmission queue names, with method of resolution
- The names of source queues with respect to:
  - Resolved destination queue names
  - Resolved destination queue manager names
  - Transmission queues
  - Message channel names
  - Adjacent system names
  - Reply-to queue names

**Note:** The use of the term *source* in this context refers to the queue name or the queue manager name provided by the application, or a channel program when opening a queue for putting messages.

An example of each of these tables is shown in Table 126, Table 127, and Table 128 on page 822.

The names in these tables are derived from the examples in this section, and this table is not intended as a practical example of queue name resolution in one node.

Table 126. Queue name resolution at queue manager QMA

| Source queue specified when queue is opened | Source queue manager specified when queue is opened | Resolved queue name | Resolved queue manager name | Resolved transmission queue name | Resolution type     |
|---------------------------------------------|-----------------------------------------------------|---------------------|-----------------------------|----------------------------------|---------------------|
| QA_norm                                     | -                                                   | QA_norm             | QMB                         | QMB                              | Remote queue        |
| (any)                                       | QMB                                                 | -                   | -                           | QMB                              | (none)              |
| QA_norm                                     | -                                                   | QA_norm             | QMB                         | TX1                              | Remote queue        |
| QB                                          | QMC                                                 | QB                  | QMD                         | QMB                              | Queue manager alias |

Table 127. Queue name resolution at queue manager QMB

| Source queue specified when queue is opened | Source queue manager specified when queue is opened | Resolved queue name | Resolved queue manager name | Resolved transmission queue name | Resolution type     |
|---------------------------------------------|-----------------------------------------------------|---------------------|-----------------------------|----------------------------------|---------------------|
| QA_norm                                     | -                                                   | QA_norm             | QMB                         | -                                | (none)              |
| QA_norm                                     | QMB                                                 | QA_norm             | QMB                         | -                                | (none)              |
| QA_norm                                     | QMB_PRIORITY                                        | QA_norm             | QMB                         | -                                | Queue manager alias |
| (any)                                       | QMC                                                 | (any)               | QMC                         | QMC                              | (none)              |
| (any)                                       | QMD_norm                                            | (any)               | QMD_norm                    | TX1                              | Queue manager alias |
| (any)                                       | QMD_PRIORITY                                        | (any)               | QMD_PRIORITY                | QMD_fast                         | Queue manager alias |

Table 127. Queue name resolution at queue manager QMB (continued)

| Source queue specified when queue is opened | Source queue manager specified when queue is opened | Resolved queue name | Resolved queue manager name | Resolved transmission queue name | Resolution type     |
|---------------------------------------------|-----------------------------------------------------|---------------------|-----------------------------|----------------------------------|---------------------|
| (any)                                       | QMC_small                                           | (any)               | QMC_small                   | TX_small                         | Queue manager alias |
| (any)                                       | QMC_large                                           | (any)               | QMC_large                   | TX_external                      | Queue manager alias |
| QB_small                                    | QMC                                                 | QB_small            | QMC                         | TX_small                         | Remote queue        |
| QB_large                                    | QMC                                                 | QB_large            | QMC                         | TX_large                         | Remote queue        |
| (any)                                       | QME                                                 | (any)               | QME                         | TX1                              | Queue manager alias |
| QA                                          | QMC                                                 | QA                  | QMC                         | TX1                              | Remote queue        |
| QB                                          | QMD                                                 | QB                  | QMD                         | TX1                              | Remote queue        |

Table 128. Reply-to queue name translation at queue manager QMA

| Application design |                                | Reply-to alias definition       |                                   |
|--------------------|--------------------------------|---------------------------------|-----------------------------------|
| Local QMGR<br>QMA  | Queue name for messages<br>QRR | Reply-to queue alias name<br>QR | Redefined to<br>QRR at QMA_class1 |

## Channel message sequence numbering

The channel uses sequence numbers to assure that messages are delivered, delivered without duplication, and stored in the same order as they were taken from the transmission queue.

The sequence number is generated at the sending end of the channel and is incremented by one before being used, which means that the current sequence number is the number of the last message sent. This information can be displayed using DISPLAY CHSTATUS. The sequence number and an identifier called the LUWID are stored in persistent storage for the last message transferred in a batch. These values are used during channel start-up to ensure that both ends of the link agree on which messages have been transferred successfully.

## Sequential retrieval of messages

If an application puts a sequence of messages to the same destination queue, those messages can be retrieved in sequence by a *single* application with a sequence of MQGET operations, if the following conditions are met:

- All the put requests were done from the same application.
- All the put requests were either from the same unit of work, or all the put requests were made outside of a unit of work.
- The messages all have the same priority.
- The messages all have the same persistence.
- For remote queuing, the configuration is such that there can only be one path from the application making the put request, through its queue manager, through intercommunication, to the destination queue manager and the target queue.
- The messages are not put to a dead-letter queue (for example, if a queue is temporarily full).
- The application getting the message does not deliberately change the order of retrieval, for example by specifying a particular *MsgId* or *CorrelId* or by using message priorities.
- Only one application is doing get operations to retrieve the messages from the destination queue. If there is more than one application, these applications must be designed to get all the messages in each sequence put by a sending application.

**Note:** Messages from other tasks and units of work might be interspersed with the sequence, even where the sequence was put from within a single unit of work.

If these conditions cannot be met, and the order of messages on the target queue is important, then the application can be coded to use its own message sequence number as part of the message to assure the order of the messages.

### **Sequence of retrieval of fast, nonpersistent messages**

Nonpersistent messages on a fast channel might overtake persistent messages on the same channel and so arrive out of sequence. The receiving MCA puts the nonpersistent messages on the destination queue immediately and makes them visible. Persistent messages are not made visible until the next sync point.

### **Loopback testing**

*Loopback testing* is a technique on non- z/OS platforms that allows you to test a communications link without actually linking to another machine.

You set up a connection between two queue managers as though they are on separate machines, but you test the connection by looping back to another process on the same machine. This technique means that you can test your communications code without requiring an active network.

The way you do so depends on which products and protocols you are using.

On Windows systems, you can use the "loopback" adapter.

Refer to the documentation for the products you are using for more information.

### **Route tracing and activity recording**

You can confirm the route a message takes through a series of queue managers in two ways.

You can use the IBM MQ display route application, available through the control command `dspmqrte`, or you can use activity recording. Both of these topics are described in Monitoring reference.

## **Introduction to distributed queue management**

Distributed queue management (DQM) is used to define and control communication between queue managers.

Distributed queue management:






- Enables you to define and control communication channels between queue managers
- Provides you with a message channel service to move messages from a type of *local queue*, known as a transmission queue, to communication links on a local system, and from communication links to local queues at a destination queue manager
- Provides you with facilities for monitoring the operation of channels and diagnosing problems, using panels, commands, and programs

Channel definitions associate channel names with transmission queues, communication link identifiers, and channel attributes. Channel definitions are implemented in different ways on different platforms. Message sending and receiving is controlled by programs known as *message channel agents* (MCAs), which use the channel definitions to start up and control communication.







The MCAs in turn are controlled by DQM itself. The structure is platform-dependent, but typically includes listeners and trigger monitors, together with operator commands and panels.

A *message channel* is a one-way pipe for moving messages from one queue manager to another. Thus a message channel has two end-points, represented by a pair of MCAs. Each end point has a definition of its end of the message channel. For example, one end would define a sender, the other end a receiver.

For details of how to define channels, see:

-    “Monitoring and controlling channels on Windows, UNIX and Linux platforms” on page 856
-  “Monitoring and controlling channels on z/OS” on page 1264
-  “Monitoring and controlling channels on IBM i” on page 879

For message channel planning examples, see:

-    Message channel planning example for distributed platforms
-  Message channel planning example for IBM MQ for IBM i
-  Message channel planning example for z/OS
-  Message channel planning example for z/OS using queue-sharing groups

For information about channel exits, see Channel-exit programs for messaging channels.

#### **Related concepts:**

“Message sending and receiving”

The following figure shows the distributed queue management model, detailing the relationships between entities when messages are transmitted. It also shows the flow for control.

“Channel control function” on page 832

The channel control function provides facilities for you to define, monitor, and control channels.

“What happens when a message cannot be delivered?” on page 847

When a message cannot be delivered, the MCA can process it in several ways. It can try again, it can return-to-sender, or it can put it on the dead-letter queue.

“Initialization and configuration files” on page 852

The handling of channel initialization data depends on your IBM MQ platform.

“Data conversion” on page 853

IBM MQ messages might require data conversion when sent between queues on different queue managers.

“Writing your own message channel agents” on page 853

IBM MQ allows you to write your own message channel agent (MCA) programs or to install one from an independent software vendor.

“Other things to consider for distributed queue management” on page 854

Other topics to consider when preparing IBM MQ for distributed queue management. This topic covers Undelivered-message queue, Queues in use, System extensions and user-exit programs, and Running channels and listeners as trusted applications.

#### **Related information:**

Example configuration information

### **Message sending and receiving**

The following figure shows the distributed queue management model, detailing the relationships between entities when messages are transmitted. It also shows the flow for control.



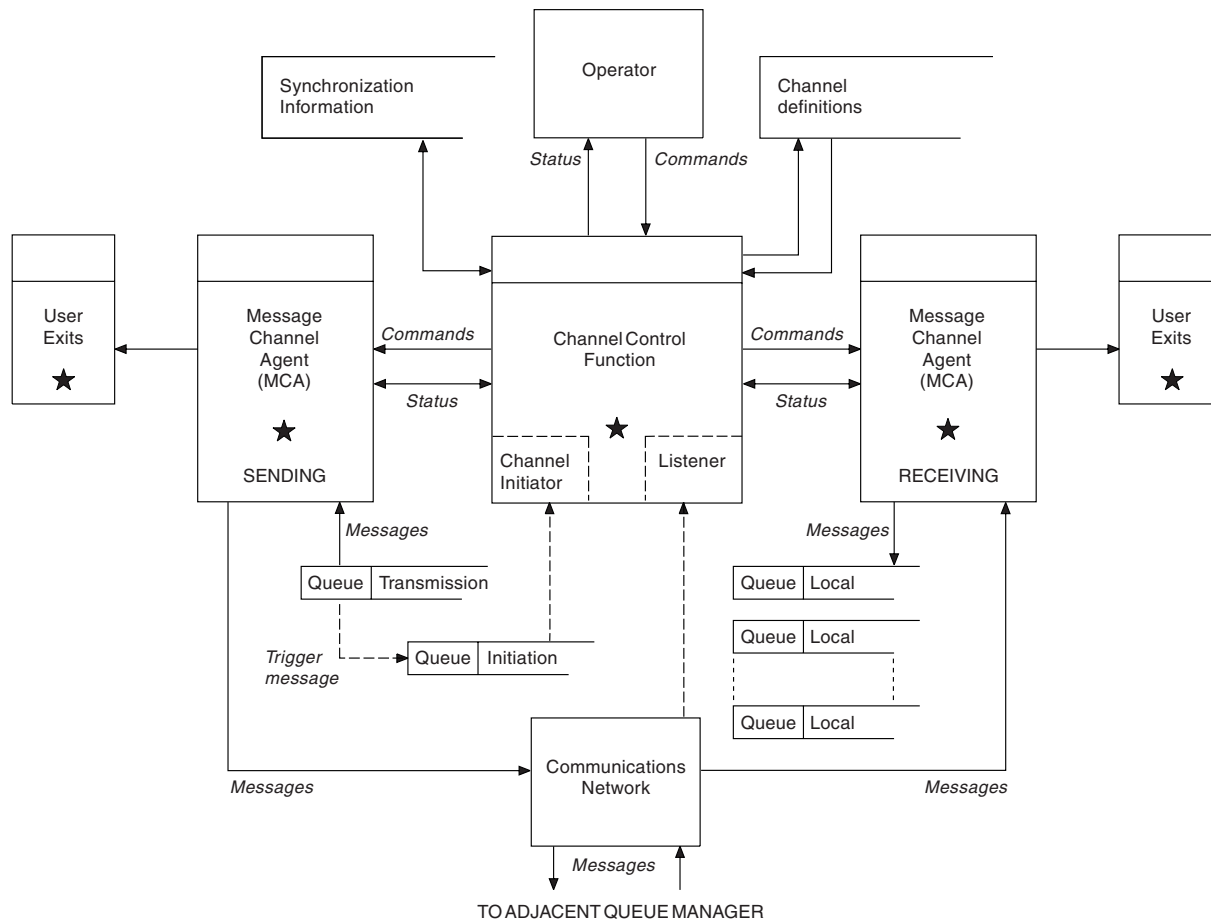


Figure 95. Distributed queue management model

**Note:**

1. There is one MCA per channel, depending on the platform. There might be one or more channel control functions for a particular queue manager.
2. The implementation of MCAs and channel control functions is highly platform-dependent. They can be programs or processes or threads, and they can be a single entity or many comprising several independent or linked parts.
3. All components marked with a star can use the MQI.

**Channel parameters**

An MCA receives its parameters in one of several ways:

- If started by a command, the channel name is passed in a data area. The MCA then reads the channel definition directly to obtain its attributes.
- For sender, and in some cases server channels, the MCA can be started automatically by the queue manager trigger. The channel name is retrieved from the trigger process definition, where applicable, and is passed to the MCA. The remaining processing is the same as previously described. Server channels must only be set up to trigger if they are fully qualified, that is, they specify a CONNAME to connect to.
- If started remotely by a sender, server, requester, or client-connection, the channel name is passed in the initial data from the partner message channel agent. The MCA reads the channel definition directly to obtain its attributes.

Certain attributes not defined in the channel definition are also negotiable:

### Split messages

If one end does not support split messages then the split messages are not sent.

### Conversion capability

If one end cannot perform the necessary code page conversion or numeric encoding conversion when needed, the other end must handle it. If neither end supports it, when needed, the channel cannot start.

### Distribution list support

If one end does not support distribution lists, the partner MCA sets a flag in its transmission queue so that it knows to intercept messages intended for multiple destinations.

## Channel status and sequence numbers

Message channel agent programs keep records of the current sequence number and logical unit of work number for each channel, and of the general status of the channel. Some platforms allow you to display this status information to help you control channels.

## How to send a message to another queue manager


This section describes the simplest way to send a message between queue managers, including prerequisites and authorizations required. Other methods can also be used to send messages to a remote queue manager.

Before you send a message from one queue manager to another, you need to do the following steps:



1. Check that your chosen communication protocol is available.
2. Start the queue managers.
3. Start the channel initiators.
4. Start the listeners.

You also need to have the correct IBM MQ security authorization to create the objects required.

To send messages from one queue manager to another:

- Define the following objects on the source queue manager:
  - Sender channel
  - Remote queue definition
  - Initiation queue (  required on z/OS, otherwise optional)
  - Transmission queue
  - Dead-letter queue
- Define the following objects on the target queue manager:
  - Receiver channel
  - Target queue
  - Dead-letter queue

You can use several different methods to define these objects, depending on your IBM MQ platform:

- On all platforms, you can use the IBM MQ script commands (MQSC) described in *The MQSC commands*, the programmable command format (PCF) commands described in *Automating administration tasks*, or the IBM MQ Explorer.
-  On z/OS, you can also use the Operation and Control panels described in *Administering IBM MQ for z/OS*.
-  On IBM i, you can also use the panel interface.

See the following subtopics for more information on creating the components for sending messages to another queue manager:

**Related concepts:**

“Creating and managing queue managers on distributed platforms” on page 687

Before you can use messages and queues, you must create and start at least one queue manager and its associated objects.

“IBM MQ distributed queuing techniques” on page 802

The subtopics in this section describe techniques that are of use when planning channels. These subtopics describe techniques to help you plan how to connect your queue managers together, and manage the flow of messages between your applications.

“Introduction to distributed queue management” on page 823

Distributed queue management (DQM) is used to define and control communication between queue managers.

“Triggering channels” on page 848

IBM MQ provides a facility for starting an application automatically when certain conditions on a queue are met. This facility is called triggering.

“Safety of messages” on page 846

In addition to the typical recovery features of IBM MQ, distributed queue management ensures that messages are delivered properly by using a sync point procedure coordinated between the two ends of the message channel. If this procedure detects an error, it closes the channel so that you can investigate the problem, and keeps the messages safely in the transmission queue until the channel is restarted.

“Monitoring and controlling channels on Windows, UNIX and Linux platforms” on page 856


For DQM you need to create, monitor, and control the channels to remote queue managers. You can control channels using commands, programs, IBM MQ Explorer, files for the channel definitions, and a storage area for synchronization information.

 “Monitoring and controlling channels on IBM i” on page 879

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers. Each queue manager has a DQM program for controlling interconnections to compatible remote queue managers.

“Configuring connections between the server and client” on page 696

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

 “Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

**Related tasks:**

“Configuring a queue manager cluster” on page 902

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

**Defining the channels:**

To send messages from one queue manager to another, you must define two channels. You must define one channel on the source queue manager and one channel on the target queue manager.

**On the source queue manager**

Define a channel with a channel type of SENDER. You need to specify the following:

- The name of the transmission queue to be used (the XMITQ attribute).
- The connection name of the partner system (the CONNAME attribute).

- The name of the communication protocol you are using (the TRPTYPE attribute). On IBM MQ for z/OS, the protocol must be TCP or LU6.2. On other platforms, you do not have to specify this. You can leave it to pick up the value from your default channel definition.

Details of all the channel attributes are given in Channel attributes.

### On the target queue manager

Define a channel with a channel type of RECEIVER, and the same name as the sender channel.

Specify the name of the communication protocol you are using (the TRPTYPE attribute). On IBM MQ for z/OS, the protocol must be TCP or LU6.2. On other platforms, you do not have to specify this. You can leave it to pick up the value from your default channel definition.

Receiver channel definitions can be generic. This means that if you have several queue managers communicating with the same receiver, the sending channels can all specify the same name for the receiver, and one receiver definition applies to them all.

When you have defined the channel, you can test it using the PING CHANNEL command. This command sends a special message from the sender channel to the receiver channel and checks that it is returned.

### Defining the queues:

To send messages from one queue manager to another, you must define up to six queues. You must define up to four queues on the source queue manager, and up to two queues on the target queue manager.

#### On the source queue manager

- Remote queue definition

In this definition, specify the following:

##### Remote queue manager name

The name of the target queue manager.


##### Remote queue name



The name of the target queue on the target queue manager.

##### Transmission queue name

The name of the transmission queue. You do not have to specify this transmission queue name. If you do not, a transmission queue with the same name as the target queue manager is used. If this does not exist, the default transmission queue is used. You are advised to give the transmission queue the same name as the target queue manager so that the queue is found by default.

- Initiation queue definition

 This is required. You must use the initiation queue called SYSTEM.CHANNEL.INITQ.

  This is optional. Consider naming the initiation queue SYSTEM.CHANNEL.INITQ.

- Transmission queue definition

A local queue with the USAGE attribute set to XMITQ.  If you are using the IBM MQ for IBM i native interface, the USAGE attribute is \*TMQ.

- Dead-letter queue definition

Define a dead-letter queue to which undelivered messages can be written.

#### On the target queue manager

- Local queue definition

The target queue. The name of this queue must be the same as that specified in the remote queue name field of the remote queue definition on the source queue manager.

- Dead-letter queue definition

Define a dead-letter queue to which undelivered messages can be written.

#### **Related concepts:**

“Creating a transmission queue”

Before a channel (other than a requester channel) can be started, the transmission queue must be defined as described in this section. The transmission queue must be named in the channel definition.

 “Creating a transmission queue on IBM i”

You can create a transmission queue on the IBM i platform by using the Create MQM Queue panel.

*Creating a transmission queue:*

Before a channel (other than a requester channel) can be started, the transmission queue must be defined as described in this section. The transmission queue must be named in the channel definition.

Define a local queue with the USAGE attribute set to XMITQ for each sending message channel. If you want to use a specific transmission queue in your remote queue definitions, create a remote queue as shown.

To create a transmission queue, use the IBM MQ Commands (MQSC), as shown in the following examples:

#### **Create transmission queue example**

```
DEFINE QLOCAL(QM2) DESCR('Transmission queue to QM2') USAGE(XMITQ)
```

#### **Create remote queue example**

```
DEFINE QREMOTE(PAYROLL) DESCR('Remote queue for QM2') +
XMITQ(QM2) RNAME(PAYROLL) RQMNAME(QM2)
```

Consider naming the transmission queue the queue manager name on the remote system, as shown in the examples.

*Creating a transmission queue on IBM i:*

You can create a transmission queue on the IBM i platform by using the Create MQM Queue panel.

You must define a local queue with the Usage field attribute set to \*TMQ, for each sending message channel.

If you want to use remote queue definitions, use the same command to create a queue of type \*RMT, and Usage of \*NORMAL.

To create a transmission queue, use the CRTMQMQ command from the command line to present you with the first queue creation panel; see Figure 96 on page 830.

```

Create MQM Queue (CRTMQMQ)
Type choices, press Enter.
Queue name
Queue type ____ *ALS, *LCL, *MDL, *RMT
Message Queue Manager name . . . *DFT_____

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
+

```

Figure 96. Create a queue (1)

Type the name of the queue and specify the type of queue that you want to create: Local, Remote, or Alias. For a transmission queue, specify Local ( \*LCL) on this panel and press enter.

You are presented with the second page of the Create MQM Queue panel; see Figure 97.

```

Create MQM Queue (CRTMQMQ)
Type choices, press Enter.
Queue name > HURS.2.HURS.PRIORIT
Queue type > *LCL *ALS, *LCL, *MDL, *RMT
Message Queue Manager name . . . *DFT
Replace *NO *NO, *YES
Text 'description' ' '
Put enabled *YES *SYSDFTQ, *NO, *YES
Default message priority 0 0-9, *SYSDFTQ
Default message persistence . . *NO *SYSDFTQ, *NO, *YES
Process name ' '
Triggering enabled *NO *SYSDFTQ, *NO, *YES
Get enabled *YES *SYSDFTQ, *NO, *YES
Sharing enabled *YES *SYSDFTQ, *NO, *YES

More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 97. Create a queue (2)

Change any of the default values shown. Press page down to scroll to the next screen; see Figure 98 on page 831.

```

Create MQM Queue (CRTMQMQ)

Type choices, press Enter.

Default share option *YES *SYSDFTQ, *NO, *YES
Message delivery sequence . . . *PTY *SYSDFTQ, *PTY, *FIFO
Harden backout count *NO *SYSDFTQ, *NO, *YES
Trigger type *FIRST *SYSDFTQ, *FIRST, *ALL...
Trigger depth 1 1-999999999, *SYSDFTQ
Trigger message priority 0 0-9, *SYSDFTQ
Trigger data ' '
Retention interval 999999999 0-999999999, *SYSDFTQ
Maximum queue depth 5000 1-24000, *SYSDFTQ
Maximum message length 4194304 0-4194304, *SYSDFTQ
Backout threshold 0 0-999999999, *SYSDFTQ
Backout requeue queue ' '
Initiation queue ' '

More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 98. Create a queue (3)

Type \*TMQ, for transmission queue, in the Usage field of this panel, and change any of the default values shown in the other fields.

```

Create MQM Queue (CRTMQMQ)

Type choices, press Enter.

Usage *TMQ *SYSDFTQ, *NORMAL, *TMQ
Queue depth high threshold . . . 80 0-100, *SYSDFTQ
Queue depth low threshold . . . 20 0-100, *SYSDFTQ
Queue full events enabled . . . *YES *SYSDFTQ, *NO, *YES
Queue high events enabled . . . *YES *SYSDFTQ, *NO, *YES
Queue low events enabled . . . *YES *SYSDFTQ, *NO, *YES
Service interval 999999999 0-999999999, *SYSDFTQ
Service interval events . . . *NONE *SYSDFTQ, *HIGH, *OK, *NONE
Distribution list support . . . *NO *SYSDFTQ, *NO, *YES
Cluster Name *SYSDFTQ
Cluster Name List *SYSDFTQ
Default Binding *SYSDFTQ *SYSDFTQ, *OPEN, *NOTFIXED

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 99. Create a queue (4)

When you are satisfied that the fields contain the correct data, press enter to create the queue.

## Starting the channel:

When you put messages on the remote queue defined at the source queue manager, they are stored on the transmission queue until the channel is started. When the channel has been started, the messages are delivered to the target queue on the remote queue manager.

Start the channel on the sending queue manager using the START CHANNEL command. When you start the sending channel, the receiving channel is started automatically (by the listener) and the messages are sent to the target queue. Both ends of the message channel must be running for messages to be transferred.

Because the two ends of the channel are on different queue managers, they could have been defined with different attributes. To resolve any differences, there is an initial data negotiation between the two ends when the channel starts. In general, the two ends of the channel operate with the attributes needing the fewer resources. This enables larger systems to accommodate the lesser resources of smaller systems at the other end of the message channel.

The sending MCA splits large messages before sending them across the channel. They are reassembled at the remote queue manager. This is not apparent to the user.

An MCA can transfer messages using multiple threads. This process, called *pipelining* enables the MCA to transfer messages more efficiently, with fewer wait states. Pipelining improves channel performance.

## Channel control function

The channel control function provides facilities for you to define, monitor, and control channels.

Commands are issued through panels, programs, or from a command line to the channel control function. The panel interface also displays channel status and channel definition data. You can use Programmable Command Formats or those IBM MQ commands (MQSC) and control commands that are detailed in "Monitoring and controlling channels on Windows, UNIX and Linux platforms" on page 856.

The commands fall into the following groups:

- Channel administration
- Channel control
- Channel status monitoring

Channel administration commands deal with the definitions of the channels. They enable you to:

- Create a channel definition
- Copy a channel definition
- Alter a channel definition
- Delete a channel definition

Channel control commands manage the operation of the channels. They enable you to:

- Start a channel
- Stop a channel
- Re-synchronize with partner (in some implementations)
- Reset message sequence numbers
- Resolve an in-doubt batch of messages
- Ping; send a test communication across the channel

Channel monitoring displays the state of channels, for example:

- Current channel settings



- Whether the channel is active or inactive
- Whether the channel terminated in a synchronized state

For more information about defining, controlling and monitoring channels, see the following subtopics:

### **Preparing channels:**

Before trying to start a message channel or MQI channel, you must prepare the channel. You must make sure that all the attributes of the local and remote channel definitions are correct and compatible.

Channel attributes describes the channel definitions and attributes.

Although you set up explicit channel definitions, the channel negotiations carried out when a channel starts, might override one or other of the values defined. This behavior is normal, and not apparent to the user, and has been arranged in this way so that otherwise incompatible definitions can work together.

### **Auto-definition of receiver and server-connection channels**

In IBM MQ on all platforms except z/OS, if there is no appropriate channel definition, then for a receiver or server-connection channel that has auto-definition enabled, a definition is created automatically. The definition is created using:

1. The appropriate model channel definition, SYSTEM.AUTO.RECEIVER, or SYSTEM.AUTO.SVRCONN. The model channel definitions for auto-definition are the same as the system defaults, SYSTEM.DEF.RECEIVER, and SYSTEM.DEF.SVRCONN, except for the description field, which is "Auto-defined by" followed by 49 blanks. The systems administrator can choose to change any part of the supplied model channel definitions.
2. Information from the partner system. The values from the partner are used for the channel name and the sequence number wrap value.
3. A channel exit program, which you can use to alter the values created by the auto-definition. See Channel auto-definition exit program.

The description is then checked to determine whether it has been altered by an auto-definition exit or because the model definition has been changed. If the first 44 characters are still "Auto-defined by" followed by 29 blanks, the queue manager name is added. If the final 20 characters are still all blanks the local time and date are added.

When the definition has been created and stored the channel start proceeds as though the definition had always existed. The batch size, transmission size, and message size are negotiated with the partner.

### **Defining other objects**

Before a message channel can be started, both ends must be defined (or enabled for auto-definition) at their queue managers. The transmission queue it is to serve must be defined to the queue manager at the sending end. The communication link must be defined and available. It might be necessary for you to prepare other IBM MQ objects, such as remote queue definitions, queue manager alias definitions, and reply-to queue alias definitions, to implement the scenarios described in "Configuring distributed queuing" on page 802.

For information about defining MQI channels, see "Defining MQI channels" on page 712.

### **Multiple message channels per transmission queue**

It is possible to define more than one channel per transmission queue, but only one of these channels can be active at any one time. Consider this option for the provision of alternative routes between queue managers for traffic balancing and link failure corrective action. A transmission queue cannot be used by

another channel if the previous channel to use it terminated leaving a batch of messages in-doubt at the sending end. For more information, see "In-doubt channels" on page 844.

### Starting a channel

A channel can be caused to start transmitting messages in one of four ways. It can be:

- Started by an operator (not receiver, cluster-receiver, or server-connection channels).
- Triggered from the transmission queue. This method applies to sender channels and fully qualified server channels (those channels which specify a CONNAME) only. You must prepare the necessary objects for triggering channels.
- Started from an application program (not receiver, cluster-receiver, or server-connection channels).
- Started remotely from the network by a sender, cluster-sender, requester, server, or client-connection channel. Receiver, cluster-receiver, and possibly server and requester channel transmissions, are started this way; so are server-connection channels. The channels themselves must already be started (that is, enabled).

**Note:** Because a channel is 'started' it is not necessarily transmitting messages. Instead, it might be 'enabled' to start transmitting when one of the four events previously described occurs. The enabling and disabling of a channel is achieved using the START and STOP operator commands.

### Channel states:

A channel can be in one of many states at any time. Some states also have substates. From a given state a channel can move into other states.

Figure 100 on page 835 shows the hierarchy of all possible channel states and the substates that apply to each of the channel states.

Figure 101 on page 836 shows the links between channel states. These links apply to all types of message channel and server-connection channels.

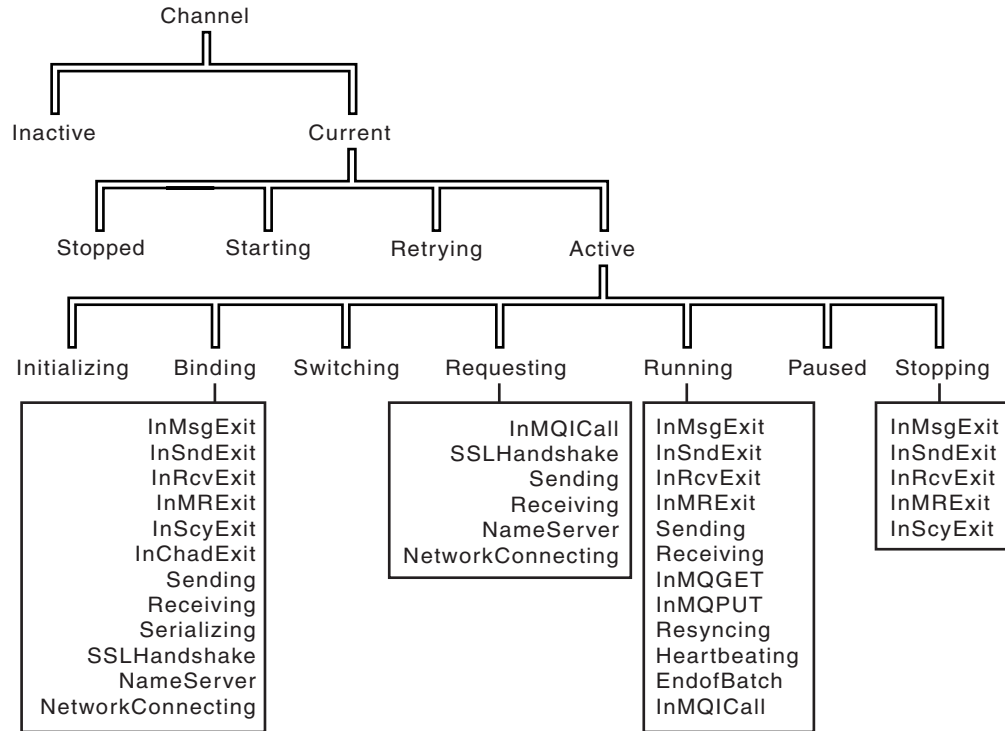


Figure 100. Channel states and substates

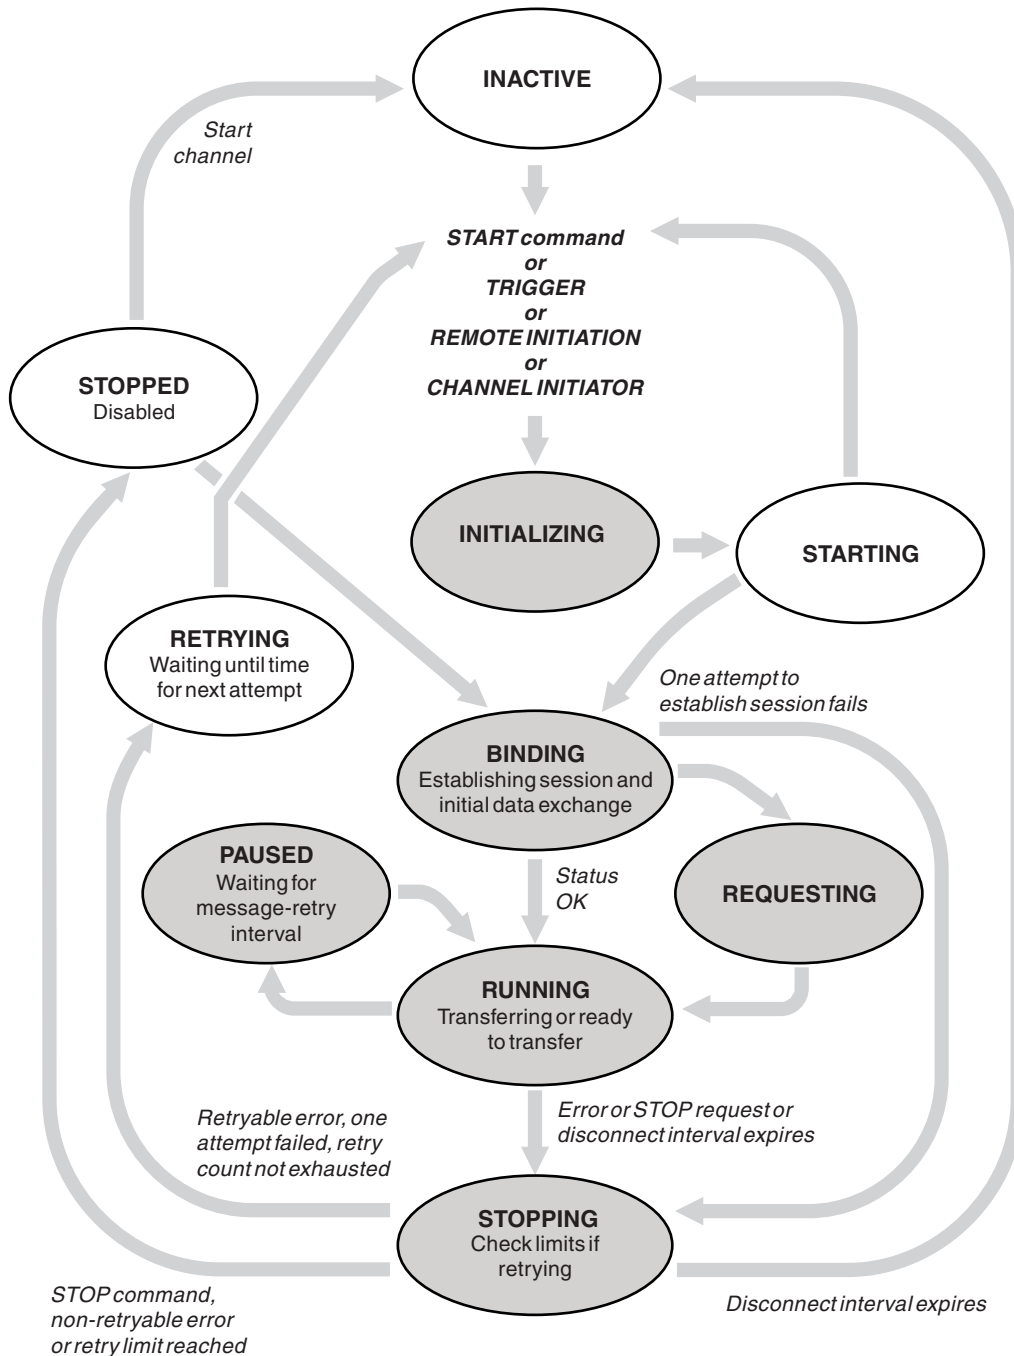


Figure 101. Flows between channel states

### Current and active

A channel is *current* if it is in any state other than inactive. A current channel is *active* unless it is in RETRYING, STOPPED, or STARTING state. When a channel is active, it is consuming resource and a process or thread is running. The seven possible states of an active channel (INITIALIZING, BINDING, SWITCHING, REQUESTING, RUNNING, PAUSED, or STOPPING) are highlighted in Figure 101.

An active channel can also show a substate giving more detail of exactly what the channel is doing. The substates for each state are shown in Figure 100 on page 835.

Current and active:

The channel is "current" if it is in any state other than inactive. A current channel is "active" unless it is in RE TRYING, STOPPED, or STARTING state.

If a channel is "active" it might also show a substate giving more detail of exactly what the channel is doing.

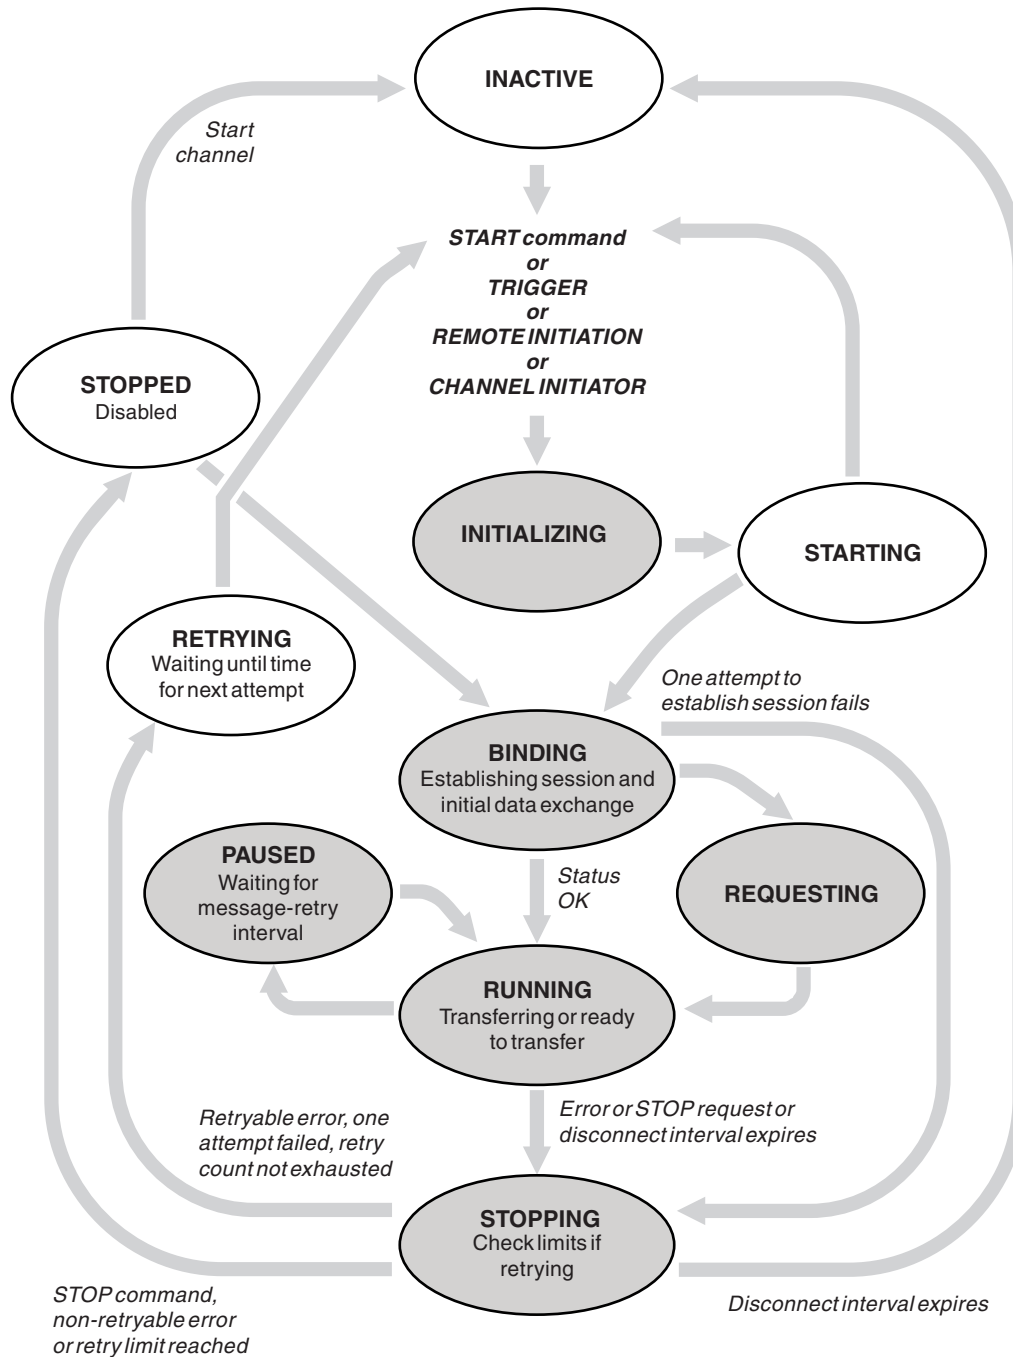






Figure 102. Flows between channel states

Note:

1. When a channel is in one of the six states highlighted in Figure 102 on page 837 (INITIALIZING, BINDING, REQUESTING, RUNNING, PAUSED, or STOPPING), it is consuming resource and a process or thread is running; the channel is *active*.
2. When a channel is in STOPPED state, the session might be active because the next state is not yet known.

### Specifying the maximum number of current channels

You can specify the maximum number of channels that can be current at one time. This number is the number of channels that have entries in the channel status table, including channels that are retrying and channels that are stopped. Specify this using  **ALTER QMGR MAXCHL** for z/OS,

 the queue manager initialization file for IBM i, the queue manager configuration file for UNIX and Linux systems, or the MQ Explorer. For more information about the values set using the initialization or the configuration file see Configuration file stanzas for distributed queuing. For more information about specifying the maximum number of channels, see Administering IBM MQ for IBM MQ for UNIX and Linux systems, and Windows systems , Administering IBM i for IBM MQ for IBM i , or Administering IBM MQ for z/OS for IBM MQ for z/OS .

#### Note:

1. Server-connection channels are included in this number.
2. A channel must be current before it can become active. If a channel is started, but cannot become current, the start fails.

### Specifying the maximum number of active channels


You can also specify the maximum number of active channels to prevent your system being overloaded by many starting channels. If you use this method, set the disconnect interval attribute to a low value to allow waiting channels to start as soon as other channels terminate.

Each time a channel that is retrying attempts to establish connection with its partner, it must become an active channel. If the attempt fails, it remains a current channel that is not active, until it is time for the next attempt. The number of times that a channel retries, and how often, is determined by the retry count and retry interval channel attributes. There are short and long values for both these attributes. See Channel attributes for more information.

When a channel has to become an active channel (because a START command has been issued, or because it has been triggered, or because it is time for another retry attempt), but is unable to do so because the number of active channels is already at the maximum value, the channel waits until one of the active slots is freed by another channel instance ceasing to be active. If, however, a channel is starting because it is being initiated remotely, and there are no active slots available for it at that time, the remote initiation is rejected.



Whenever a channel, other than a requester channel, is attempting to become active, it goes into the STARTING state. This state occurs even if there is an active slot immediately available, although it is only in the STARTING state for a short time. However, if the channel has to wait for an active slot, it is in STARTING state while it is waiting.

Requester channels do not go into STARTING state. If a requester channel cannot start because the number of active channels is already at the limit, the channel ends abnormally.

Whenever a channel, other than a requester channel, is unable to get an active slot, and so waits for one, a message is written to the log  or the z/OS console, and an event is generated. When a slot is later freed and the channel is able to acquire it, another message and event are generated. Neither of these events and messages are generated if the channel is able to acquire a slot straight away.

If a STOP CHANNEL command is issued while the channel is waiting to become active, the channel goes to STOPPED state. A Channel-Stopped event is raised.


Server-connection channels are included in the maximum number of active channels.

For more information about specifying the maximum number of active channels, see *Administering IBM MQ for IBM MQ for UNIX and Linux systems, and Windows systems* , *Administering IBM i for IBM MQ for IBM i* , or *Administering IBM MQ for z/OS for IBM MQ for z/OS*.


#### *Channel errors:*



Errors on channels cause the channel to stop further transmissions. If the channel is a sender or server, it goes to RETRY state because it is possible that the problem might clear itself. If it cannot go to RETRY state, the channel goes to STOPPED state.

For sending channels, the associated transmission queue is set to GET(DISABLED) and triggering is turned off. (A STOP command with STATUS(STOPPED) takes the side that issued it to STOPPED state; only expiry of the disconnect interval or a STOP command with STATUS(INACTIVE) makes it end normally and become inactive.) Channels that are in STOPPED state need operator intervention before they can restart (see “Restarting stopped channels” on page 844 ).

**Note:** For  IBM i, UNIX, Linux and Windows systems, a channel initiator must be running for retry to be attempted. If the channel initiator is not available, the channel becomes inactive and must be manually restarted. If you are using a script to start the channel, ensure that the channel initiator is running before you try to run the script.

Long retry count (LONGRTY) describes how retrying works. If the error clears, the channel restarts automatically, and the transmission queue is re-enabled. If the retry limit is reached without the error clearing, the channel goes to STOPPED state. A stopped channel must be restarted manually by the operator. If the error is still present, it does not retry again. When it does start successfully, the transmission queue is re-enabled.

 If the channel initiator stops while a channel is in RETRYING or STOPPED status, the channel status is remembered when the channel initiator is restarted. However, the channel status for the SVRCONN channel type is reset if the channel initiator stops while the channel is in STOPPED status.

  If the queue manager stops while a channel is in RETRYING or STOPPED status, the channel status is remembered when the queue manager is restarted. However, the channel status for the SVRCONN channel type is reset if the queue manager stops while the channel is in STOPPED status.

If a channel is unable to put a message to the target queue because that queue is full or put inhibited, the channel can retry the operation a number of times (specified in the message-retry count attribute) at a time interval (specified in the message-retry interval attribute). Alternatively, you can write your own message-retry exit that determines which circumstances cause a retry, and the number of attempts made. The channel goes to PAUSED state while waiting for the message-retry interval to finish.

See *Channel attributes* for information about the channel attributes, and *Channel-exit programs* for messaging channels for information about the message-retry exit.

### Server-connection channel limits:

You can set server-connection channel limits to prevent client applications from exhausting queue manager channel resources, **MAXINST**, and to prevent a single client application from exhausting server-connection channel capacity, **MAXINSTC**.

A maximum total number of channels can be active at any time on an individual queue manager. The total number of server-connection channel instances are included in the maximum number of active channels.

If you do not specify the maximum number of simultaneous instances of a server-connection channel that can be started, then it is possible for a single client application, connecting to a single server-connection channel, to exhaust the maximum number of active channels that are available. When the maximum number of active channels is reached, it prevents any other channels from being started on the queue manager. To avoid this situation, you must limit the number of simultaneous instances of an individual server-connection channel that can be started, regardless of which client started them.

If the value of the limit is reduced to below the currently running number of instances of the server connection channel, even to zero, then the running channels are not affected. New instances cannot be started until sufficient existing instances have ceased to run so that the number of currently running instances is less than the value of the limit.

Also, many different client-connection channels can connect to an individual server-connection channel. The limit on the number of simultaneous instances of an individual server-connection channel that can be started, regardless of which client started them, prevents any client from exhausting the maximum active channel capacity of the queue manager. If you do not also limit the number of simultaneous instances of an individual server-connection channel that can be started from an individual client, then it is possible for a single, faulty client application to open so many connections that it exhausts the channel capacity allocated to an individual server-connection channel, and therefore prevents other clients that need to use the channel from connecting to it. To avoid this situation, you must limit the number of simultaneous instances of an individual server-connection channel that can be started from an individual client.

If the value of the individual client limit is reduced below the number of instances of the server-connection channel that are currently running from individual clients, even to zero, then the running channels are not affected. However, new instances of the server-connection channel cannot be started from an individual client that exceeds the new limit until sufficient existing instances from that client have ceased to run so that the number of currently running instances is less than the value of this parameter.

### Checking that the other end of the channel is still available:

You can use the heartbeat interval, the keep alive interval, and the receive timeout, to check that the other end of the channel is available.

#### Heartbeats

You can use the heartbeat interval channel attribute to specify that flows are to be passed from the sending MCA when there are no messages on the transmission queue, as is described in Heartbeat interval (HBINT).

#### Keep alive

In IBM MQ for z/OS, if you are using TCP/IP as the transport protocol, you can also specify a value for the **Keepalive** interval channel attribute (KAINT). You are recommended to give the **Keepalive** interval a higher value than the heartbeat interval, and a smaller value than the disconnect value. You can use this attribute to specify a time-out value for each channel, as is described in Keepalive Interval (KAINT).



In IBM MQ for IBM i, UNIX, Linux, and Windows systems, if you are using TCP as your transport protocol, you can set `keepalive=yes`. If you specify this option, TCP periodically checks that the other end of the connection is still available. If it is not, the channel is terminated. This option is described in Keepalive Interval (KAINTE).

If you have unreliable channels that report TCP errors, use of the **Keepalive** option means that your channels are more likely to recover.

You can specify time intervals to control the behavior of the **Keepalive** option. When you change the time interval, only TCP/IP channels started after the change are affected. Ensure that the value that you choose for the time interval is less than the value of the disconnect interval for the channel.

For more information about using the **Keepalive** option, see the KAINTE parameter in the DEFINE CHANNEL command.

### Receive timeout

If you are using TCP as your transport protocol, the receiving end of an idle non-MQI channel connection is also closed if no data is received for a period. This period, the *receive time-out* value, is determined according to the HBINT (heartbeat interval) value.

In IBM MQ for IBM i, UNIX, Linux, and Windows systems, the *receive time-out* value is set as follows:

1. For an initial number of flows, before any negotiation takes place, the *receive time-out* value is twice the HBINT value from the channel definition.
2. After the channels negotiate an HBINT value, if HBINT is set to less than 60 seconds, the *receive time-out* value is set to twice this value. If HBINT is set to 60 seconds or more, the *receive time-out* value is set to 60 seconds greater than the value of HBINT.

In IBM MQ for z/OS, the *receive time-out* value is set as follows:

1. For an initial number of flows, before any negotiation takes place, the *receive time-out* value is twice the HBINT value from the channel definition.
2. If RCVTIME is set, the timeout is set to one of
  - the negotiated HBINT multiplied by a constant
  - the negotiated HBINT plus a constant number of seconds
  - a constant number of seconds

depending on the RCVTTYPE parameter, and subject to any limit imposed by RCVTMIN if it applies. RCVTMIN does not apply when RCVTTYPE(EQUAL) is configured. If you use a constant value of RCVTIME and you use a heartbeat interval, do not specify an RCVTIME less than the heartbeat interval. For details of the RCVTIME, RCVTMIN and RCVTTYPE attributes, see the ALTER QMGR command.

### Note:

1. If either of the values is zero, there is no timeout.
2. For connections that do not support heartbeats, the HBINT value is negotiated to zero in step 2 and hence there is no timeout, so you must use TCP/IP KEEPALIVE.
3. For client connections that use sharing conversations, heartbeats can flow across the channel (from both ends) all the time, not just when an MQGET is outstanding.
4. For client connections where sharing conversations are not in use, heartbeats are flowed from the server only when the client issues an MQGET call with wait. Therefore, you are not recommended to set the heartbeat interval too small for client channels. For example, if the heartbeat is set to 10 seconds, an MQCMIT call fails (with MQRC\_CONNECTION\_BROKEN) if it takes longer than 20 seconds to commit because no data flowed during this time. This can happen with large units of

work. However, it does not happen if appropriate values are chosen for the heartbeat interval because only MQGET with wait takes significant periods of time.

Provided SHARECNV is not zero, the client uses a full duplex connection, which means that the client can (and does) heartbeat during all MQI calls

5. In IBM MQ Version 7 Client channels, heartbeats can flow from both the server as well as the client side. The timeout at either end is based upon  $2 \times \text{HBINT}$  for HBINTs of less than 60 seconds and  $\text{HBINT} + 60$  for HBINTs of over 60 seconds.
6. Canceling the connection after twice the heartbeat interval is valid because a data or heartbeat flow is expected at least at every heartbeat interval. Setting the heartbeat interval too small, however, can cause problems, especially if you are using channel exits. For example, if the HBINT value is one second, and a send or receive exit is used, the receiving end waits for only 2 seconds before canceling the channel. If the MCA is performing a task such as encrypting the message, this value might be too short.




### Adopting an MCA:

The Adopt MCA function enables IBM MQ to cancel a receiver channel and start a new one in its place.

If a channel loses contact, the receiver channel can be left in a 'communications receive' state. When communications are re-established the sender channel attempts to reconnect. If the remote queue manager finds that the receiver channel is already running it does not allow another version of the same receiver channel to start. This problem requires user intervention to rectify the problem or the use of system keepalive.

The Adopt MCA function solves the problem automatically. It enables IBM MQ to cancel a receiver channel and to start a new one in its place.

The function can be set up with various options:

-  For distributed platforms, see Administering IBM MQ.
-  For z/OS, see Administering IBM MQ for z/OS.
-  For IBM i, see Administering IBM i.



### Stopping and quiescing channels:

You can stop and quiesce a channel before the disconnect time interval expires.

Message channels are designed to be long-running connections between queue managers with orderly termination controlled only by the disconnect interval channel attribute. This mechanism works well unless the operator needs to terminate the channel before the disconnect time interval expires. This need can occur in the following situations:

- System quiesce
- Resource conservation
- Unilateral action at one end of a channel

In this case, you can stop the channel. You can do this using:

- the STOP CHANNEL MQSC command
- the Stop Channel PCF command
- the IBM MQ Explorer
-   other platform-specific mechanisms, as follows:

 **For z/OS:**

The Stop a channel panel

The ENDMQMCHL CL command or the END option on the WRKMQMCHL panel

There are three options for stopping channels using these commands:

### QUIESCE

The QUIESCE option attempts to end the current batch of messages before stopping the channel.

### FORCE

The FORCE option attempts to stop the channel immediately and might require the channel to resynchronize when it restarts because the channel might be left in doubt.

**z/OS** On IBM MQ for z/OS, FORCE interrupts any message reallocation in progress, which might leave BIND\_NOT\_FIXED messages partially reallocated or out of order.

### TERMINATE

The TERMINATE option attempts to stop the channel immediately, and terminates the thread or process of the channel.

**z/OS** On IBM MQ for z/OS, TERMINATE interrupts any message reallocation in progress, which might leave BIND\_NOT\_FIXED messages partially reallocated or out of order.

All these options leave the channel in a STOPPED state, requiring operator intervention to restart it.

Stopping the channel at the sending end is effective but does require operator intervention to restart. At the receiving end of the channel, things are much more difficult because the MCA is waiting for data from the sending side, and there is no way to initiate an *orderly* termination of the channel from the receiving side; the stop command is pending until the MCA returns from its wait for data.

Consequently there are three recommended ways of using channels, depending upon the operational characteristics required:

- If you want your channels to be long running, note that there can be orderly termination only from the sending end. When channels are interrupted, that is, stopped, operator intervention (a START CHANNEL command) is required in order to restart them.
- If you want your channels to be active only when there are messages for them to transmit, set the disconnect interval to a fairly low value. The default setting is high and so is not recommended for channels where this level of control is required. Because it is difficult to interrupt the receiving channel, the most economical option is to have the channel automatically disconnect and reconnect as the workload demands. For most channels, the appropriate setting of the disconnect interval can be established heuristically.
- You can use the heartbeat-interval attribute to cause the sending MCA to send a heartbeat flow to the receiving MCA during periods in which it has no messages to send. This action releases the receiving MCA from its wait state and gives it an opportunity to quiesce the channel without waiting for the disconnect interval to expire. Give the heartbeat interval a lower value than the value of the disconnect interval.

### Note:

1. You are advised to set the disconnect interval to a low value, or to use heartbeats, for server channels. This low value is to allow for the case where the requester channel ends abnormally (for example, because the channel was canceled) when there are no messages for the server channel to send. If the disconnect interval is set high and heartbeats are not in use, the server does not detect that the requester has ended (which it will only do the next time it tries to send a message to the requester). While the server is still running, it holds the transmission queue open for exclusive input in order to get any more messages that arrive on the queue. If an attempt is made to restart the channel from the requester, the start request receives an error because the server still has the transmission queue open for exclusive input. It is necessary to stop the server channel, and then restart the channel from the requester again.

## Restarting stopped channels:

When a channel goes into STOPPED state, you have to restart the channel manually.

To do restart the channel, issue one of the following commands:

- The START CHANNEL MQSC command
- The Start Channel PCF command
- the IBM MQ Explorer



other platform-specific mechanisms, as follows:


 **For z/OS:**



The Start a channel panel

 **For IBM i:**

The STRMQMCHL CL command or the START option on the WRKMQMCHL panel

For sender or server channels, when the channel entered the STOPPED state, the associated transmission queue was set to GET(DISABLED) and triggering was set off. When the start request is received, these attributes are reset automatically.

 If the channel initiator stops while a channel is in RETRYING or STOPPED status, the channel status is remembered when the channel initiator is restarted. However, the channel status for the SVRCONN channel type is reset if the channel initiator stops while the channel is in STOPPED status.

  If the queue manager stops while a channel is in RETRYING or STOPPED status, the channel status is remembered when the queue manager is restarted. From IBM MQ Version 8.0 onwards, this applies to SVRCONN channels as well. Previously, the channel status for the SVRCONN channel type was reset if the channel initiator stopped while the channel was in STOPPED status.

## In-doubt channels:

An in-doubt channel is a channel that is in doubt with a remote channel about which messages have been sent and received.

Note the distinction between this and a queue manager being in doubt about which messages should be committed to a queue.

You can reduce the opportunity for a channel to be placed in doubt by using the Batch Heartbeat channel parameter (BATCHHB). When a value for this parameter is specified, a sender channel checks that the remote channel is still active before taking any further action. If no response is received the receiver channel is considered to be no longer active. The messages can be rolled-back, and re-routed, and the sender-channel is not put in doubt. This reduces the time when the channel could be placed in doubt to the period between the sender channel verifying that the receiver channel is still active, and verifying that the receiver channel has received the sent messages. See Channel attributes for more information about the batch heartbeat parameter.

In-doubt channel problems are typically resolved automatically. Even when communication is lost, and a channel is placed in doubt with a message batch at the sender with receipt status unknown, the situation is resolved when communication is re-established. Sequence number and LUWID records are kept for this purpose. The channel is in doubt until LUWID information has been exchanged, and only one batch of messages can be in doubt for the channel.

You can, when necessary, resynchronize the channel manually. The term *manual* includes use of operators or programs that contain IBM MQ system management commands. The manual resynchronization process works as follows. This description uses MQSC commands, but you can also use the PCF equivalents.

1. Use the DISPLAY CHSTATUS command to find the last-committed logical unit of work ID (LUWID) for *each* side of the channel. Do this using the following commands:

- For the in-doubt side of the channel:

```
DISPLAY CHSTATUS(name) SAVED CURLUWID
```

You can use the CONNAME and XMITQ parameters to further identify the channel.

- For the receiving side of the channel:

```
DISPLAY CHSTATUS(name) SAVED LSTLUWID
```

You can use the CONNAME parameter to further identify the channel.

The commands are different because only the sending side of the channel can be in doubt. The receiving side is never in doubt.


On IBM MQ for IBM i, the DISPLAY CHSTATUS command can be executed from a file using the STRMQMMQSC command or the Work with MQM Channel Status CL command, WRKMQMCHST

2. If the two LUWIDs are the same, the receiving side has committed the unit of work that the sender considers to be in doubt. The sending side can now remove the in-doubt messages from the transmission queue and re-enable it. This is done with the following channel RESOLVE command:

```
RESOLVE CHANNEL(name) ACTION(COMMIT)
```

3. If the two LUWIDs are different, the receiving side has not committed the unit of work that the sender considers to be in doubt. The sending side needs to retain the in-doubt messages on the transmission queue and re-send them. This is done with the following channel RESOLVE command:

```
RESOLVE CHANNEL(name) ACTION(BACKOUT)
```

 On IBM MQ for IBM i, you can use the Resolve MQM Channel command, RSVMQMCHL.

Once this process is complete the channel is no longer in doubt. The transmission queue can now be used by another channel, if required.

### **Problem determination:**

There are two distinct aspects to problem determination - problems discovered when a command is being submitted, and problems discovered during operation of the channels.

### **Command validation**

Commands and panel data must be free from errors before they are accepted for processing. Any errors found by the validation are immediately notified to the user by error messages.

Problem diagnosis begins with the interpretation of these error messages and taking corrective action.

### **Processing problems**

Problems found during normal operation of the channels are notified to the system console or the system log. Problem diagnosis begins with the collection of all relevant information from the log, and continues with analysis to identify the problem.

Confirmation and error messages are returned to the terminal that initiated the commands, when possible.

IBM MQ produces accounting and statistical data, which you can use to identify trends in utilization and performance. On z/OS, this information is produced in the form of SMF records, see *Monitoring performance and resource usage* for details. The equivalent information about other platforms is produced as PCF records, see *Structure data types* for details.

## Messages and codes

For messages and codes to help with the primary diagnosis of the problem, see *Diagnostic messages and reason codes*.

## Safety of messages

In addition to the typical recovery features of IBM MQ, distributed queue management ensures that messages are delivered properly by using a sync point procedure coordinated between the two ends of the message channel. If this procedure detects an error, it closes the channel so that you can investigate the problem, and keeps the messages safely in the transmission queue until the channel is restarted.

The sync point procedure has an added benefit in that it attempts to recover an *in-doubt* situation when the channel starts. (*In-doubt* is the status of a unit of recovery for which a sync point has been requested but the outcome of the request is not yet known.) Also associated with this facility are the two functions:

1. Resolve with commit or backout
2. Reset the sequence number

The use of these functions occurs only in exceptional circumstances because the channel recovers automatically in most cases.

## Fast, nonpersistent messages

The nonpersistent message speed (NPMSPEED) channel attribute can be used to specify that any nonpersistent messages on the channel are to be delivered more quickly. For more information about this attribute, see *Nonpersistent message speed (NPMSPEED)*.

If a channel terminates while fast, nonpersistent messages are in transit, the messages might be lost and it is up to the application to arrange for their recovery if required.

If the receiving channel cannot put the message to its destination queue then it is placed on the dead letter queue, if one has been defined. If not, the message is discarded.

**Note:** If the other end of the channel does not support the option, the channel runs at normal speed.

## Undelivered Messages

For information about what happens when a message cannot be delivered, see “What happens when a message cannot be delivered?” on page 847.

## What happens when a message cannot be delivered?

When a message cannot be delivered, the MCA can process it in several ways. It can try again, it can return-to-sender, or it can put it on the dead-letter queue.

Figure 103 shows the processing that occurs when an MCA is unable to put a message to the destination queue. (The options shown do not apply on all platforms.)

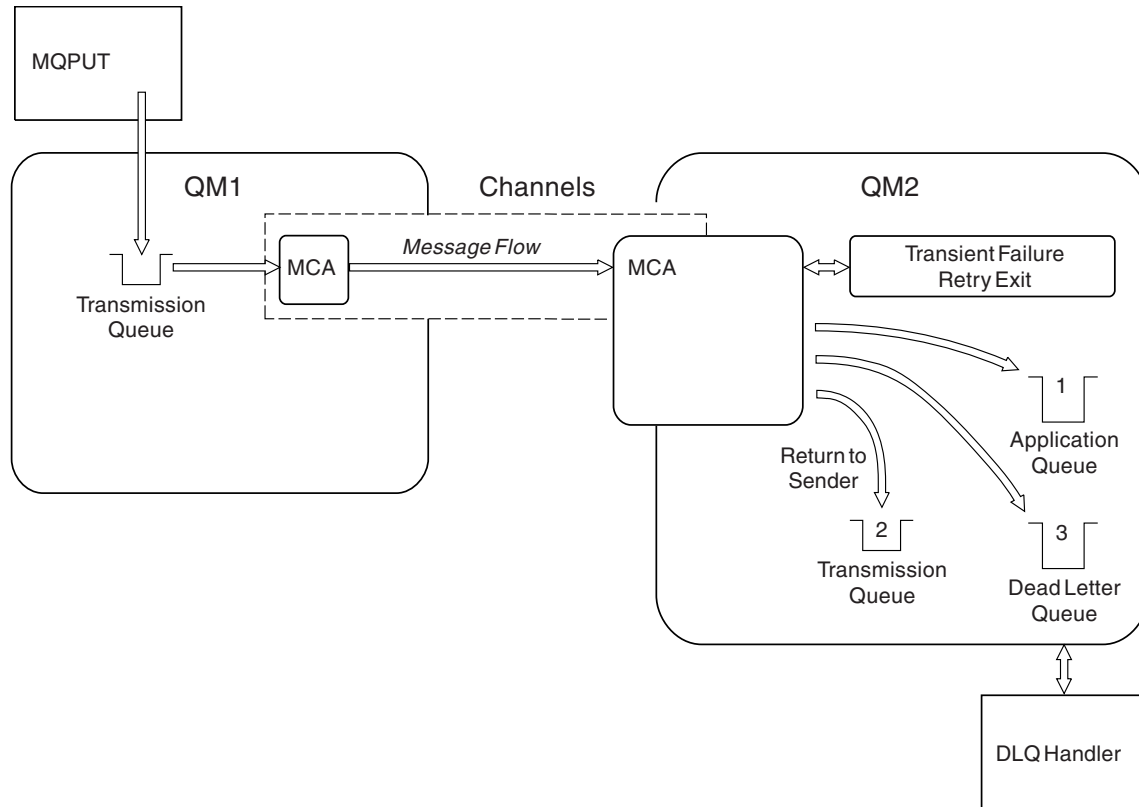


Figure 103. What happens when a message cannot be delivered

As shown in the figure, the MCA can do several things with a message that it cannot deliver. The action taken is determined by options specified when the channel is defined and on the MQPUT report options for the message.

### 1. Message-retry

If the MCA is unable to put a message to the target queue for a reason that could be transitory (for example, because the queue is full), the MCA can wait and try the operation again later. You can determine if the MCA waits, for how long, and how many times it tries.

- You can specify a message-retry time and interval for MQPUT errors when you define your channel. If the message cannot be put to the destination queue because the queue is full, or is inhibited for puts, the MCA tries the operation the number of times specified, at the time interval specified.
- You can write your own message-retry exit. The exit enables you to specify under what conditions you want the MCA to try the MQPUT or MQOPEN operation again. Specify the name of the exit when you define the channel.

### 2. Return-to-sender

If message-retry was unsuccessful, or a different type of error was encountered, the MCA can send the message back to the originator. To enable return-to-sender, you need to specify the following options in the message descriptor when you put the message to the original queue:

- The MQRO\_EXCEPTION\_WITH\_FULL\_DATA report option

- The MQRO\_DISCARD\_MSG report option
- The name of the reply-to queue and reply-to queue manager

If the MCA is unable to put the message to the destination queue, it generates an exception report containing the original message, and puts it on a transmission queue to be sent to the reply-to queue specified in the original message. (If the reply-to queue is on the same queue manager as the MCA, the message is put directly to that queue, not to a transmission queue.)

### 3. Dead-letter queue

If a message cannot be delivered or returned, it is put on to the dead-letter queue (DLQ). You can use the DLQ handler to process the message. This processing is described in Processing messages on a dead-letter queue for IBM MQ for UNIX, Linux and Windows systems, and in The dead-letter queue handler utility (CSQUDLQH) for z/OS systems. If the dead-letter queue is not available, the sending MCA leaves the message on the transmission queue, and the channel stops. On a fast channel, nonpersistent messages that cannot be written to a dead-letter queue are lost.

On IBM WebSphere MQ Version 7.0, if no local dead-letter queue is defined, the remote queue is not available or defined, and there is no remote dead-letter queue, then the sender channel goes into RETRY and messages are automatically rolled back to the transmission queue.

#### Related information:



Use Dead-Letter Queue (USEDLQ)

## Triggering channels

IBM MQ provides a facility for starting an application automatically when certain conditions on a queue are met. This facility is called triggering.

This explanation is intended as an overview of triggering concepts. For a complete description, see Starting IBM MQ applications using triggers.

For platform-specific information see the following:

- For Windows, see UNIX and Linux systems, “Triggering channels on UNIX, Linux and Windows systems.” on page 850
-  For IBM i, see “Triggering channels in IBM MQ for IBM i” on page 850
-  For z/OS, see “Transmission queues and triggering channels” on page 1263






Figure 104. The concepts of triggering

The objects required for triggering are shown in Figure 104. It shows the following sequence of events:

1. The local queue manager places a message from an application or from a message channel agent (MCA) on the transmission queue.
2. When the triggering conditions are fulfilled, the local queue manager places a trigger message on the initiation queue.
3. The long-running channel initiator program monitors the initiation queue, and retrieves messages as they arrive.
4. The channel initiator processes the trigger messages according to information contained in them. This information might include the channel name, in which case the corresponding MCA is started.
5. The local application or the MCA, having been triggered, retrieves the messages from the transmission queue.

To set up this scenario, you need to:

- Create the transmission queue with the name of the initiation queue (that is, `SYSTEM.CHANNEL.INITQ`) in the corresponding attribute.
- Ensure that the initiation queue (`SYSTEM.CHANNEL.INITQ`) exists.
- Ensure that the channel initiator program is available and running. The channel initiator program must be provided with the name of the initiation queue in its start command. z/OS On z/OS, the name of the initiation queue is fixed, so is not used on the start command.
- Optionally, create the process definition for the triggering, if it does not exist, and ensure that the *UserData* field contains the name of the channel it serves. Instead of creating a process definition, you can specify the channel name in the *TriggerData* attribute of the transmission queue. IBM MQ for IBM i IBM i, UNIX, Linux and Windows systems, allow the channel name to be specified as blank, in which case the first available channel definition with this transmission queue is used.
- Ensure that the transmission queue definition contains the name of the process definition to serve it, (if applicable), the initiation queue name, and the triggering characteristics you feel are most suitable. The trigger control attribute allows triggering to be enabled, or not, as necessary.

**Note:**

1. The channel initiator program acts as a 'trigger monitor' monitoring the initiation queue used to start channels.
2. An initiation queue and trigger process can be used to trigger any number of channels.
3. Any number of initiation queues and trigger processes can be defined.
4. A trigger type of FIRST is recommended, to avoid flooding the system with channel starts.

### Triggering channels on UNIX, Linux and Windows systems.

You can create a process definition in IBM MQ, defining processes to be triggered. Use the MQSC command DEFINE PROCESS to create a process definition naming the process to be triggered when messages arrive on a transmission queue. The USERDATA attribute of the process definition contains the name of the channel being served by the transmission queue.

Define the local queue (QM4), specifying that trigger messages are to be written to the initiation queue (IQ) to trigger the application that starts channel (QM3.TO.QM4):

```
DEFINE QLOCAL(QM4) TRIGGER INITQ(SYSTEM.CHANNEL.INITQ) PROCESS(P1) USAGE(XMITQ)
```

Define the application (process P1) to be started:

```
DEFINE PROCESS(P1) USERDATA(QM3.TO.QM4)
```

Alternatively, for IBM MQ for UNIX, Linux and Windows systems, you can eliminate the need for a process definition by specifying the channel name in the TRIGDATA attribute of the transmission queue.

Define the local queue (QM4). Specify that trigger messages are to be written to the default initiation queue SYSTEM.CHANNEL.INITQ, to trigger the application (process P1) that starts channel (QM3.TO.QM4):

```
DEFINE QLOCAL(QM4) TRIGGER INITQ(SYSTEM.CHANNEL.INITQ)
USAGE(XMITQ) TRIGDATA(QM3.TO.QM4)
```

If you do not specify a channel name, the channel initiator searches the channel definition files until it finds a channel that is associated with the named transmission queue.



### Triggering channels in IBM MQ for IBM i

Triggering of channels in IBM MQ for IBM i is implemented with the channel initiator process. A channel initiator process for the initiation queue SYSTEM.CHANNEL.INITQ is started automatically with the queue manager unless it is disabled by altering the queue manager SCHINIT attribute.

Set up the transmission queue for the channel, specifying SYSTEM.CHANNEL.INITQ as the initiation queue, and enabling triggering for the queue. The channel initiator starts the first available channel that specifies this transmission queue.

```
CRTMQMQ QNAME(MYXMITQ1) QTYPE(*LCL) MQMNAME(MYQMGR)
TRGENBL(*YES) INITQNAME(SYSTEM.CHANNEL.INITQ)
USAGE(*TMQ)
```

You can manually start up to three channel initiator processes with the STRMQMCHLI command and specify different initiation queues. You can also specify more than one channel able to process the transmission queue and choose which channel to start. This capability is still provided to be compatible with earlier releases. Its usage is deprecated.

**Note:** Only one channel at a time can process a transmission queue.

```
STRMQMCHLI QNAME(MYINITQ)
```

Set up the transmission queue for the channel, specifying TRGENBL(\*YES) and, to choose which channel to attempt to start, specify the channel name in the TRIGDATA field. For example:

```
CRTMQMQ QNAME(MYXMITQ2) QTYPE(*LCL) MQMNAME(MYQMGR)
TRGENBL(*YES) INITQNAME(MYINITQ)
USAGE(*TMQ) TRIGDATA(MYCHANNEL)
```

#### Related concepts:

“Starting and stopping the channel initiator”


Triggering is implemented using the channel initiator process.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

#### Related information:

Channel programs on UNIX, Linux, and Windows systems

 Intercommunication jobs on IBM i

 Channel states on IBM i

#### Starting and stopping the channel initiator:


Triggering is implemented using the channel initiator process.

This channel initiator process is started with the MQSC command START CHINIT. Unless you are using the default initiation queue, specify the name of the initiation queue on the command. For example, to use the START CHINIT command to start queue IQ for the default queue manager, enter:

```
START CHINIT INITQ(IQ)
```

By default, a channel initiator is started automatically using the default initiation queue, SYSTEM.CHANNEL.INITQ. If you want to start all your channel initiators manually, follow these steps:

1. Create and start the queue manager.
2. Alter the queue manager's SCHINIT property to MANUAL
3. End and restart the queue manager

In  IBM MQ for iSeries, UNIX, Linux and Windows systems, a channel initiator is started automatically. The number of channel initiators that you can start is limited. The default and maximum value is 3. You can change this using MAXINITIATORS in the qm.ini file for UNIX and Linux systems, and in the registry for Windows systems.

See IBM MQ Control commands for details of the run channel initiator command **runmqchi**, and the other control commands.

#### Stopping the channel initiator

The default channel initiator is started automatically when you start a queue manager. All channel initiators are stopped automatically when a queue manager is stopped.

## Initialization and configuration files

The handling of channel initialization data depends on your IBM MQ platform.

▶ z/OS

### z/OS systems

In IBM MQ for z/OS, initialization and configuration information is specified using the ALTER QMGR MQSC command. If you put ALTER QMGR commands in the CSQINP2 initialization input data set, they are processed every time the queue manager is started.

To run MQSC commands such as START LISTENER every time you start the channel initiator, put them in the CSQINPX initialization input data set and specify the optional DD statement CSQINPX in the channel initiator started task procedure.

For further information about CSQINP2 and CSQINPX, see *Customize the initialization input data sets, and ALTER QMGR*.

### Windows, IBM i, IBM i, UNIX and Linux systems

In IBM MQ for Windows, IBM i, IBM i, UNIX and Linux systems, there are *configuration files* to hold basic configuration information about the IBM MQ installation.

There are two configuration files: one applies to the machine, the other applies to an individual queue manager.

#### IBM MQ configuration file

This file holds information relevant to all the queue managers on the IBM MQ system. The file is called `mqs.ini`. It is fully described in the *Administering for IBM MQ for Windows, IBM i, and in UNIX and Linux systems*.

#### Queue manager configuration file

This file holds configuration information relating to one particular queue manager. The file is called `qm.ini`.

It is created during queue manager creation and can hold configuration information relevant to any aspect of the queue manager. Information held in the file includes details of how the configuration of the log differs from the default in the IBM MQ configuration file.

The queue manager configuration file is held in the root of the directory tree occupied by the queue manager. For example, for the `DefaultPath` attributes, the queue manager configuration files for a queue manager called `QMNAME` would be:

For UNIX and Linux systems:

```
/var/mqm/qmgrs/QMNAME/qm.ini
```

An excerpt of a `qm.ini` file follows. It specifies that the TCP/IP listener is to listen on port 2500, the maximum number of current channels is to be 200, and the maximum number of active channels is to be 100.


```
TCP:
Port=2500
CHANNELS:
MaxChannels=200
MaxActiveChannels=100
```

You can specify a range of TCP/IP ports to be used by an outbound channel. One method is to use the `qm.ini` file to specify the start and end of a range of port values. The following example shows a `qm.ini` file specifying a range of channels:

```
TCP:
StrPort=2500
EndPort=3000
CHANNELS:
MaxChannels=200
MaxActiveChannels=100
```

If you specify a value for StrPort or EndPort then you must specify a value for both. The value of EndPort must always be greater than the value of StrPort.

The channel tries to use each of the port values in the range specified. When the connection is successful, the port value is the port that the channel then uses.

 For IBM i:  
/QIBM/UserData/mqm/qmgrs/QMNAME/qm.ini

For Windows systems:  
C:\ProgramData\IBM\MQ\qmgrs\QMNAME\qm.ini

For more information about qm.ini files, see Configuration file stanzas for distributed queuing.

## Data conversion

IBM MQ messages might require data conversion when sent between queues on different queue managers.

An IBM MQ message consists of two parts:

- Control information in a message descriptor
- Application data

Either of the two parts might require data conversion when sent between queues on different queue managers. For information about application data conversion, see Application data conversion.

## Writing your own message channel agents

IBM MQ allows you to write your own message channel agent (MCA) programs or to install one from an independent software vendor.

You might want to write your own MCA programs to make IBM MQ interoperate over your own proprietary communications protocol, or to send messages over a protocol that IBM MQ does not support. (You cannot write your own MCA to interoperate with an IBM MQ-supplied MCA at the other end.)

If you decide to use an MCA that was not supplied by IBM MQ, you must consider the following points.

### Message sending and receiving

You must write a sending application that gets messages from wherever your application puts them, for example from a transmission queue, and sends them out on a protocol with which you want to communicate. You must also write a receiving application that takes messages from this protocol and puts them onto destination queues. The sending and receiving applications use the message queue interface (MQI) calls, not any special interfaces.

You must ensure that messages are only delivered once. Sync point coordination can be used to help with this delivery.

### Channel control function

You must provide your own administration functions to control channels. You cannot use IBM MQ channel administration functions either for configuring (for example, the DEFINE CHANNEL command) or monitoring (for example, DISPLAY CHSTATUS) your channels.

### Initialization file

You must provide your own initialization file, if you require one.

### Application data conversion

You probably want to allow for data conversion for messages you send to a different system. If so, use the MQGMO\_CONVERT option on the MQGET call when retrieving messages from wherever your application puts them, for example the transmission queue.

### User exits

Consider whether you need user exits. If so, you can use the same interface definitions that IBM MQ uses.

### Triggering

If your application puts messages to a transmission queue, you can set up the transmission queue attributes so that your sending MCA is triggered when messages arrive on the queue.


### Channel initiator

You might must provide your own channel initiator.

## Other things to consider for distributed queue management

Other topics to consider when preparing IBM MQ for distributed queue management. This topic covers Undelivered-message queue, Queues in use, System extensions and user-exit programs, and Running channels and listeners as trusted applications.


### Undelivered-message queue

To ensure that messages arriving on the undelivered-message queue (also known as the dead-letter queue or DLQ) are processed, create a program that can be triggered or run at regular intervals to handle these messages. A DLQ handler is provided with IBM MQ on UNIX and Linux systems; for more information, see The sample DLQ handler, amqsdlq.  For more information on IBM MQ for IBM i, see The IBM MQ for IBM i dead-letter queue handler.

### Queues in use

MCAs for receiver channels can keep the destination queues open even when messages are not being transmitted. This results in the queues appearing to be “in use”.

### Maximum number of channels

 On IBM MQ for IBM i you can specify the maximum number of channels allowed in your system and the maximum number that can be active at one time. You specify these numbers in the qm.ini file in directory QIBM/UserData/mqm/qmgrs/*queue\_manager\_name*. See Configuration file stanzas for distributed queuing.

### System extensions and user-exit programs

A facility is provided in the channel definition to enable extra programs to be run at defined times during the processing of messages. These programs are not supplied with IBM MQ, but can be provided by each installation according to local requirements.

In order to run, these user-exit programs must have predefined names and be available on call to the channel programs. The names of the user-exit programs are included in the message channel definitions.

There is a defined control block interface for handing over control to these programs, and for handling the return of control from these programs.

The precise places where these programs are called, and details of control blocks and names, are to be found in Channel-exit programs for messaging channels.

## Running channels and listeners as trusted applications

If performance is an important consideration in your environment and your environment is stable, you can run your channels and listeners as trusted, using the FASTPATH binding. There are two factors that influence whether channels and listeners run as trusted:

- The environment variable `MQ_CONNECT_TYPE=FASTPATH` or `MQ_CONNECT_TYPE=STANDARD`. This is case-sensitive. If you specify a value that is not valid it is ignored.
- `MQIBindType` in the Channels stanza of the `qm.ini` or registry file. You can set this to `FASTPATH` or `STANDARD` and it is not case-sensitive. The default is `STANDARD`.

You can use `MQIBindType` in association with the environment variable to achieve the required effect as follows:

| <b>MQIBindType</b> | <b>Environment variable</b> | <b>Result</b> |
|--------------------|-----------------------------|---------------|
| STANDARD           | UNDEFINED                   | STANDARD      |
| FASTPATH           | UNDEFINED                   | FASTPATH      |
| STANDARD           | STANDARD                    | STANDARD      |
| FASTPATH           | STANDARD                    | STANDARD      |
| STANDARD           | FASTPATH                    | STANDARD      |
| FASTPATH           | FASTPATH                    | FASTPATH      |
| STANDARD           | CLIENT                      | CLIENT        |
| FASTPATH           | CLIENT                      | STANDARD      |
| STANDARD           | LOCAL                       | STANDARD      |
| FASTPATH           | LOCAL                       | STANDARD      |

In summary, there are only two ways of actually making channels and listeners run as trusted:

1. By specifying `MQIBindType=FASTPATH` in `qm.ini` or registry and not specifying the environment variable.
2. By specifying `MQIBindType=FASTPATH` in `qm.ini` or registry and setting the environment variable to `FASTPATH`.

Consider running listeners as trusted, because listeners are stable processes. Consider running channels as trusted, unless you are using unstable channel exits or the command `STOP CHANNEL MODE(TERMINATE)`.

## Monitoring and controlling channels on Windows, UNIX and Linux platforms

For DQM you need to create, monitor, and control the channels to remote queue managers. You can control channels using commands, programs, IBM MQ Explorer, files for the channel definitions, and a storage area for synchronization information.

You can use the following types of command:

### The IBM MQ commands (MQSC)

You can use the MQSC as single commands in an MQSC session in Windows, UNIX and Linux systems. To issue more complicated, or multiple, commands the MQSC can be built into a file that you then run from the command line. For details, see MQSC commands. This section gives some simple examples of using MQSC for distributed queuing.

The channel commands are a subset of the IBM MQ Commands (MQSC). You use MQSC and the control commands to:

- Create, copy, display, change, and delete channel definitions
- Start and stop channels, ping, reset channel sequence numbers, and resolve in-doubt messages when links cannot be re-established
- Display status information about channels

### Control commands

You can also issue *control commands* at the command line for some of these functions. For details, see control commands.

### Programmable command format commands

For details, see PCF commands.

### IBM MQ Explorer

On UNIX, Linux and Windows systems, you can use the IBM MQ Explorer. This provides a graphical administration interface to perform administrative tasks as an alternative to using control commands or MQSC commands. Channel definitions are held as queue manager objects.

Each queue manager has a DQM component for controlling interconnections to compatible remote queue managers. A storage area holds sequence numbers and *logical unit of work (LUW)* identifiers. These are used for channel synchronization purposes.

For a list of the functions available to you when setting up and controlling message channels, using the different types of command, see Table 129 on page 857.

### Related concepts:

“Getting started with objects” on page 859

Channels must be defined, and their associated objects must exist and be available for use, before a channel can be started. This section shows you how.

“Setting up communication for Windows” on page 865

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. For this to succeed, it is necessary for the connection to be defined and available. This section explains how to do this using one of the four forms of communication for IBM MQ for Windows systems.

“Setting up communication on UNIX and Linux systems” on page 873

DQM is a remote queuing facility for IBM MQ. It provides channel control programs for the queue manager which form the interface to communication links, controllable by the system operator. The channel definitions held by distributed-queuing management use these connections.

 “Monitoring and controlling channels on IBM i” on page 879

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers. Each queue manager has a DQM program for controlling interconnections to compatible



remote queue managers.

**Related information:**

- Channel programs on UNIX, Linux, and Windows systems
- Message channel planning example for distributed platforms
- Example configuration information
- Channel attributes

## Functions required for setting up and controlling channels

A number of IBM MQ functions might be needed to set up and control channels. The channel functions are explained in this topic.

You can create a channel definition using the default values supplied by IBM MQ, specifying the name of the channel, the type of channel you are creating, the communication method to be used, the transmission queue name and the connection name.

The channel name must be the same at both ends of the channel, and unique within the network. However, you must restrict the characters used to those that are valid for IBM MQ object names.

For other channel related functions, see the following topics:

- “Getting started with objects” on page 859
- “Creating associated objects” on page 859
- “Creating default objects” on page 859
- “Creating a channel” on page 860
- “Displaying a channel” on page 861
- “Displaying channel status” on page 861
- “Checking links using Ping” on page 862
- “Starting a channel” on page 862
- “Stopping a channel” on page 863
- “Renaming a channel” on page 864
- “Resetting a channel” on page 864
- “Resolving in-doubt messages on a channel” on page 865

Table 129 shows the full list of IBM MQ functions that you might need.

*Table 129. Functions required in Windows, UNIX and Linux systems*

| Function                 | Control commands | MQSC         | IBM MQ Explorer equivalent? |
|--------------------------|------------------|--------------|-----------------------------|
| Queue manager functions  |                  |              |                             |
| Change queue manager     |                  | ALTER QMGR   | Yes                         |
| Create queue manager     | crtmqm           |              | Yes                         |
| Delete queue manager     | dltmqm           |              | Yes                         |
| Display queue manager    |                  | DISPLAY QMGR | Yes                         |
| End queue manager        | endmqm           |              | Yes                         |
| Ping queue manager       |                  | PING QMGR    | No                          |
| Start queue manager      | strmqm           |              | Yes                         |
| Command server functions |                  |              |                             |
| Display command server   | dspmqcsv         |              | No                          |
| End command server       | endmqcsv         |              | No                          |

Table 129. Functions required in Windows, UNIX and Linux systems (continued)

| Function                  | Control commands                                          | MQSC                                                                                         | IBM MQ Explorer equivalent? |
|---------------------------|-----------------------------------------------------------|----------------------------------------------------------------------------------------------|-----------------------------|
| Start command server      | strmqcsv                                                  |                                                                                              | No                          |
| Queue functions           |                                                           |                                                                                              |                             |
| Change queue              |                                                           | ALTER QALIAS<br>ALTER QLOCAL<br>ALTER QMODEL<br>ALTER QREMOTE<br><br>See, ALTER queues.      | Yes                         |
| Clear queue               |                                                           | CLEAR QLOCAL                                                                                 | Yes                         |
| Create queue              |                                                           | DEFINE QALIAS<br>DEFINE QLOCAL<br>DEFINE QMODEL<br>DEFINE QREMOTE<br><br>See, DEFINE queues. | Yes                         |
| Delete queue              |                                                           | DELETE QALIAS<br>DELETE QLOCAL<br>DELETE QMODEL<br>DELETE QREMOTE<br><br>See, DELETE queues. | Yes                         |
| Display queue             |                                                           | DISPLAY QUEUE                                                                                | Yes                         |
| Process functions         |                                                           |                                                                                              |                             |
| Change process            |                                                           | ALTER PROCESS                                                                                | Yes                         |
| Create process            |                                                           | DEFINE PROCESS                                                                               | Yes                         |
| Delete process            |                                                           | DELETE PROCESS                                                                               | Yes                         |
| Display process           |                                                           | DISPLAY PROCESS                                                                              | Yes                         |
| Channel functions         |                                                           |                                                                                              |                             |
| Change channel            |                                                           | ALTER CHANNEL                                                                                | Yes                         |
| Create channel            |                                                           | DEFINE CHANNEL                                                                               | Yes                         |
| Delete channel            |                                                           | DELETE CHANNEL                                                                               | Yes                         |
| Display channel           |                                                           | DISPLAY CHANNEL                                                                              | Yes                         |
| Display channel status    |                                                           | DISPLAY CHSTATUS                                                                             | Yes                         |
| End channel               |                                                           | STOP CHANNEL                                                                                 | Yes                         |
| Ping channel              |                                                           | PING CHANNEL                                                                                 | Yes                         |
| Reset channel             |                                                           | RESET CHANNEL                                                                                | Yes                         |
| Resolve channel           |                                                           | RESOLVE CHANNEL                                                                              | Yes                         |
| Run channel               | runmqchl                                                  | START CHANNEL                                                                                | Yes                         |
| Run channel initiator     | runmqchi                                                  | START CHINIT                                                                                 | No                          |
| Run listener <sup>1</sup> | runmqlsr                                                  | START LISTENER                                                                               | No                          |
| End listener              | endmqlsr ( Windows systems, AIX, HP-UX, and Solaris only) |                                                                                              | No                          |

**Note:**

1. A listener might be started automatically when the queue manager starts.

## Getting started with objects

Channels must be defined, and their associated objects must exist and be available for use, before a channel can be started. This section shows you how.

Use the IBM MQ commands (MQSC) or the IBM MQ Explorer to:

1. Define message channels and associated objects
2. Monitor and control message channels

The associated objects you might need to define are:

- Transmission queues
- Remote queue definitions
- Queue manager alias definitions
- Reply-to queue alias definitions
- Reply-to local queues
- Processes for triggering (MCAs)
- Message channel definitions

The particular communication link for each channel must be defined and available before a channel can be run. For a description of how LU 6.2, TCP/IP, NetBIOS, SPX, and DECnet links are defined, see the particular communication guide for your installation. See also Example configuration information.

For more information about creating and working with objects, see the following subtopics:

### Creating associated objects:

MQSC is used to create associated objects.

Use MQSC to create the queue and alias objects: transmission queues, remote queue definitions, queue manager alias definitions, reply-to queue alias definitions, and reply-to local queues.

Also create the definitions of processes for triggering (MCAs) in a similar way.

For an example showing how to create all the required objects see Message channel planning example for distributed platforms.

### Creating default objects:

Default objects are created automatically when a queue manager is created. These objects are queues, channels, a process definition, and administration queues. After the default objects have been created, you can replace them at any time by running the `strmqm` command with the `-c` option.

When you use the `crtmqm` command to create a queue manager, the command also initiates a program to create a set of default objects.

1. Each default object is created in turn. The program keeps a count of how many objects are successfully defined, how many existed and were replaced, and how many unsuccessful attempts there were.
2. The program displays the results to you and if any errors occurred, directs you to the appropriate error log for details.

When the program has finished running, you can use the `strmqm` command to start the queue manager.

See The control commands for more information about the `crtmqm` and `strmqm` commands.

### Changing the default objects

When you specify the `-c` option, the queue manager is started temporarily while the objects are created and is then shut down again. Issuing `strmqm` with the `-c` option refreshes existing system objects with the default values (for example, the `MCAUSER` attribute of a channel definition is set to blanks). You must use the `strmqm` command again, without the `-c` option, if you want to start the queue manager.

If you want to change the default objects, you can create your own version of the old `amqscoma.tst` file and edit it.

### Creating a channel:

Create *two* channel definitions, one at each end of the connection. You create the first channel definition at the first queue manager. Then you create the second channel definition at the second queue manager, on the other end of the link.

Both ends must be defined using the *same* channel name. The two ends must have **compatible** channel types, for example: Sender and Receiver.

To create a channel definition for one end of the link use the MQSC command `DEFINE CHANNEL`. Include the name of the channel, the channel type for this end of the connection, a connection name, a description (if required), the name of the transmission queue (if required), and the transmission protocol. Also include any other attributes that you want to be different from the system default values for the required channel type, using the information you have gathered previously.

You are provided with help in deciding on the values of the channel attributes in `Channel attributes`.

**Note:** You are recommended to name all the channels in your network uniquely. Including the source and target queue manager names in the channel name is a good way to do this.

### Create channel example

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(SDR) +
DESCR('Sender channel to QM2') +
CONNNAME(QM2) TRPTYPE(TCP) XMITQ(QM2) CONVERT(YES)
```

In all the examples of MQSC the command is shown as it appears in a file of commands, and as it is typed in Windows or UNIX or Linux systems. The two methods look identical, except that to issue a command interactively, you must first start an MQSC session. Type `runmqsc`, for the default queue manager, or `runmqsc qmname` where *qmname* is the name of the required queue manager. Then type any number of commands, as shown in the examples.

For portability, restrict the line length of your commands to 72 characters. Use the concatenation character, `+`, as shown to continue over more than one line. On Windows use `Ctrl-z` to end the entry at the command line. On UNIX and Linux systems, use `Ctrl-d`. Alternatively, on UNIX, Linux or Windows systems, use the **end** command.

### **Displaying a channel:**

Use the MQSC command `DISPLAY CHANNEL` to display the attributes of a channel.

The `ALL` parameter of the `DISPLAY CHANNEL` command is assumed by default if no specific attributes are requested and the channel name specified is not generic.

The attributes are described in [Channel attributes](#).

### **Display channel examples**

```
DISPLAY CHANNEL(QM1.TO.QM2) TRPTYPE,CONVERT
```

```
DISPLAY CHANNEL(QM1.TO.*) TRPTYPE,CONVERT
```

```
DISPLAY CHANNEL(*) TRPTYPE,CONVERT
```

```
DISPLAY CHANNEL(QM1.TO.QMR34) ALL
```

### **Displaying channel status:**

Use the MQSC command `DISPLAY CHSTATUS`, specifying the channel name and whether you want the current status of channels or the status of saved information.

`DISPLAY CHSTATUS` applies to all message channels. It does not apply to MQI channels other than server-connection channels.

Information displayed includes:

- Channel name
- Communication connection name
- In-doubt status of channel (where appropriate)
- Last sequence number
- Transmission queue name (where appropriate)
- The in-doubt identifier (where appropriate)
- The last committed sequence number
- Logical unit of work identifier
- Process ID
- Thread ID ( Windows only)

### **Display channel status examples**

```
DISPLAY CHSTATUS(*) CURRENT
```

```
DISPLAY CHSTATUS(QM1.TO.*) SAVED
```

The saved status does not apply until at least one batch of messages has been transmitted on the channel. Status is also saved when a channel is stopped (using the `STOP CHL` command) and when the queue manager is ended.

## Checking links using Ping:

Use the MQSC command `PING CHANNEL` to exchange a fixed data message with the remote end.

Ping gives some confidence to the system supervisor that the link is available and functioning.

Ping does not involve the use of transmission queues and target queues. It uses channel definitions, the related communication link, and the network setup. It can only be used if the channel is not currently active.

It is available from sender, server, and cluster-sender channels only. The corresponding channel is started at the far side of the link, and performs the startup parameter negotiation. Errors are notified normally.

The result of the message exchange is presented as Ping complete or an error message.

## Ping with LU 6.2

When Ping is invoked, by default no user ID or password flows to the receiving end. If user ID and password are required, they can be created at the initiating end in the channel definition. If a password is entered into the channel definition, it is encrypted by IBM MQ before being saved. It is then decrypted before flowing across the conversation.

## Starting a channel:

Use the MQSC command `START CHANNEL` for sender, server, and requester channels. For applications to be able to exchange messages, you must start a listener program for inbound connections.

`START CHANNEL` is not necessary where a channel has been set up with queue manager triggering.

When started, the sending MCA reads the channel definitions and opens the transmission queue. A channel start-up sequence is issued, which remotely starts the corresponding MCA of the receiver or server channel. When they have been started, the sender and server processes await messages arriving on the transmission queue and transmit them as they arrive.

When you use triggering or run channels as threads, ensure that the channel initiator is available to monitor the initiation queue. The channel initiator is started by default as part of the queue manager.

However, TCP and LU 6.2 do provide other capabilities:

- For TCP on UNIX and Linux systems, `inetd` can be configured to start a channel. `inetd` is started as a separate process.
- For LU 6.2 in UNIX and Linux systems, configure your SNA product to start the LU 6.2 responder process.
- For LU 6.2 in Windows systems, using SNA Server you can use `TpStart` (a utility provided with SNA Server) to start a channel. `TpStart` is started as a separate process.

Use of the `Start` option always causes the channel to resynchronize, where necessary.

For the start to succeed:

- Channel definitions, local and remote, must exist. If there is no appropriate channel definition for a receiver or server-connection channel, a default one is created automatically if the channel is auto-defined. See Channel auto-definition exit program.
- Transmission queue must exist, and have no other channels using it.
- MCAs, local and remote, must exist.
- Communication link must be available.

- Queue managers must be running, local and remote.
- Message channel must not be already running.

A message is returned to the screen confirming that the request to start a channel has been accepted. For confirmation that the start command has succeeded, check the error log, or use DISPLAY CHSTATUS. The error logs are:

#### Windows

*MQ\_INSTALLATION\_PATH*\qmgrs\qmname\errors\AMQERR01.LOG (for each queue manager called qmname)

*MQ\_INSTALLATION\_PATH*\qmgrs\@SYSTEM\errors\AMQERR01.LOG (for general errors)

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

**Note:** On Windows systems, you still also get a message in the Windows systems application event log.

#### UNIX and Linux systems

/var/mqm/qmgrs/qmname/errors/AMQERR01.LOG (for each queue manager called qmname)

/var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG (for general errors)

On Windows, UNIX and Linux systems, use the runmqslr command to start the IBM MQ listener process. By default, any inbound requests for channel attachment causes the listener process to start MCAs as threads of the amqrmppa process.

```
runmqslr -t tcp -m QM2
```

For outbound connections, you must start the channel in one of the following three ways:

1. Use the MQSC command START CHANNEL, specifying the channel name, to start the channel as a process or a thread, depending on the MCATYPE parameter. (If channels are started as threads, they are threads of a channel initiator.)

```
START CHANNEL(QM1.TO.QM2)
```

2. Use the control command runmqchl to start the channel as a process.

```
runmqchl -c QM1.TO.QM2 -m QM1
```

3. Use the channel initiator to trigger the channel.

#### Stopping a channel:

Use the MQSC command STOP CHANNEL to request the channel to stop activity. The channel does not start a new batch of messages until the operator starts the channel again.

For information about restarting stopped channels, see “Restarting stopped channels” on page 844.

This command can be issued to a channel of any type except MQCHT\_CLNTCONN.

You can select the type of stop you require:

#### Stop quiesce example

```
STOP CHANNEL(QM1.TO.QM2) MODE(QUIESCE)
```

This command requests the channel to close down in an orderly way. The current batch of messages is completed and the sync point procedure is carried out with the other end of the channel. If the channel is idle this command does not terminate a receiving channel.

#### Stop force example

```
STOP CHANNEL(QM1.TO.QM2) MODE(FORCE)
```

This option stops the channel immediately, but does not terminate the channel's thread or process. The channel does not complete processing the current batch of messages, and can, therefore, leave the channel in doubt. In general, consider using the quiesce stop option.

#### **Stop terminate example**

```
STOP CHANNEL(QM1.TO.QM2) MODE(TERMINATE)
```

This option stops the channel immediately, and terminates the channel's thread or process.

#### **Stop (quiesce) stopped example**

```
STOP CHANNEL(QM1.TO.QM2) STATUS(STOPPED)
```

This command does not specify a MODE, so defaults to MODE(QUIESCE). It requests that the channel is stopped so that it cannot be restarted automatically but must be started manually.

#### **Stop (quiesce) inactive example**

```
STOP CHANNEL(QM1.TO.QM2) STATUS(INACTIVE)
```

This command does not specify a MODE, so defaults to MODE(QUIESCE). It requests that the channel is made inactive so that it restarts automatically when required.

#### **Renaming a channel:**

Use MQSC to rename a message channel.

Use MQSC to carry out the following steps:

1. Use STOP CHANNEL to stop the channel.
2. Use DEFINE CHANNEL to create a duplicate channel definition with the new name.
3. Use DISPLAY CHANNEL to check that it has been created correctly.
4. Use DELETE CHANNEL to delete the original channel definition.

If you decide to rename a message channel, remember that a channel has *two* channel definitions, one at each end. Make sure that you rename the channel at both ends at the same time.

#### **Resetting a channel:**

Use the MQSC command RESET CHANNEL to change the message sequence number.

The RESET CHANNEL command is available for any message channel, but not for MQI channels (client-connection or server-connection). The first message starts the new sequence the next time the channel is started.

If the command is issued on a sender or server channel, it informs the other side of the change when the channel is restarted.



**Related concepts:**

“Getting started with objects” on page 859

Channels must be defined, and their associated objects must exist and be available for use, before a channel can be started. This section shows you how.

“Channel control function” on page 832

The channel control function provides facilities for you to define, monitor, and control channels.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

**Related information:**

RESET CHANNEL

**Resolving in-doubt messages on a channel:**

Use the MQSC command RESOLVE CHANNEL when messages are held in-doubt by a sender or server. For example because one end of the link has terminated, and there is no prospect of it recovering.

The RESOLVE CHANNEL command accepts one of two parameters: BACKOUT or COMMIT. Backout restores messages to the transmission queue, while Commit discards them.

The channel program does not try to establish a session with a partner. Instead, it determines the logical unit of work identifier (LUWID) which represents the in-doubt messages. It then issues, as requested, either:

- BACKOUT to restore the messages to the transmission queue; or
- COMMIT to delete the messages from the transmission queue.

For the resolution to succeed:

- The channel must be inactive
- The channel must be in doubt
- The channel type must be sender, server, or cluster-sender
- A local channel definition must exist
- The local queue manager must be running

**Related concepts:**

“Getting started with objects” on page 859

Channels must be defined, and their associated objects must exist and be available for use, before a channel can be started. This section shows you how.

“Channel control function” on page 832

The channel control function provides facilities for you to define, monitor, and control channels.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

**Related information:**

RESOLVE CHANNEL

**Setting up communication for Windows**

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. For this to succeed, it is necessary for the connection to be defined and available. This section explains how to do this using one of the four forms of communication for IBM MQ for Windows systems.

You might find it helpful to refer to Example configuration - IBM MQ for Windows .

For UNIX and Linux systems see “Setting up communication on UNIX and Linux systems” on page 873.

## Deciding on a connection

Choose from the following four forms of communication for IBM MQ for Windows systems:

- “Defining a TCP connection on Windows”
- “Defining an LU 6.2 connection on Windows” on page 868
- “Defining a NetBIOS connection on Windows” on page 870

Each channel definition must specify only one protocol as the Transmission protocol (Transport Type) attribute. One or more protocols can be used by a queue manager.

For IBM MQ clients, it might be useful to have alternative channels using different transmission protocols. For more information about IBM MQ clients, see Overview of clients.

### Related concepts:

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Monitoring and controlling channels on Windows, UNIX and Linux platforms” on page 856

For DQM you need to create, monitor, and control the channels to remote queue managers. You can control channels using commands, programs, IBM MQ Explorer, files for the channel definitions, and a storage area for synchronization information.

“Configuring connections between the server and client” on page 696

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

### Defining a TCP connection on Windows:

Define a TCP connection by configuring a channel at the sending end to specify the address of the target, and by running a listener program at the receiving end.

### Sending end

Specify the host name, or the TCP address of the target machine, in the Connection name field of the channel definition.

The port to connect to defaults to 1414. Port number 1414 is assigned by the Internet Assigned Numbers Authority to IBM MQ.

To use a port number other than the default, specify it in the connection name field of the channel object definition thus:

```
DEFINE CHANNEL('channel name') CHLTYPE(SDR) +
 TRPTYPE(TCP) +
 CONNAME('0S2R0G3(1822)') +
 XMITQ('XMitQ name') +
 REPLACE
```

where 0S2R0G3 is the DNS name of the remote queue manager and 1822 is the port required. (This must be the port that the listener at the receiving end is listening on.)

A running channel must be stopped and restarted to pick up any change to the channel object definition.

You can change the default port number by specifying it in the .ini file for IBM MQ for Windows:

TCP:  
Port=1822

**Note:** To select which TCP/IP port number to use, IBM MQ uses the first port number it finds in the following sequence:

1. The port number explicitly specified in the channel definition or command line. This number allows the default port number to be overridden for a channel.
2. The port attribute specified in the TCP stanza of the .ini file. This number allows the default port number to be overridden for a queue manager.
3. The default value of 1414. This is the number assigned to IBM MQ by the Internet Assigned Numbers Authority for both inbound and outbound connections.

For more information about the values you set using qm.ini, see Configuration file stanzas for distributed queuing.

### Receiving on TCP

To start a receiving channel program, a listener program must be started to detect incoming network requests and start the associated channel. You can use the IBM MQ listener.

Receiving channel programs are started in response to a startup request from the sending channel.

To start a receiving channel program, a listener program must be started to detect incoming network requests and start the associated channel. You can use the IBM MQ listener.

To run the Listener supplied with IBM MQ, that starts new channels as threads, use the runmq1sr command.

A basic example of using the **runmq1sr** command:

```
runmq1sr -t tcp [-m QMNAME] [-p 1822]
```

The square brackets indicate optional parameters; QMNAME is not required for the default queue manager, and the port number is not required if you are using the default (1414). The port number must not exceed 65535.

**Note:** To select which TCP/IP port number to use, IBM MQ uses the first port number it finds in the following sequence:

1. The port number explicitly specified in the channel definition or command line. This number allows the default port number to be overridden for a channel.
2. The port attribute specified in the TCP stanza of the .ini file. This number allows the default port number to be overridden for a queue manager.
3. The default value of 1414. This is the number assigned to IBM MQ by the Internet Assigned Numbers Authority for both inbound and outbound connections.

For the best performance, run the IBM MQ listener as a trusted application as described in “Running channels and listeners as trusted applications” on page 855. See Restrictions for trusted applications for information about trusted applications

### Using the TCP/IP SO\_KEEPALIVE option

If you want to use the Windows SO\_KEEPALIVE option you must add the following entry to your registry:

TCP:  
KeepAlive=yes

For more information about the SO\_KEEPALIVE option, see “Checking that the other end of the channel is still available” on page 840.

On Windows, the HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters registry value for the Windows KeepAliveTime option controls the interval that elapses before the connection is checked. The default is two hours.

### Defining an LU 6.2 connection on Windows:

SNA must be configured so that an LU 6.2 conversation can be established between the two machines.

Once the SNA is configured, proceed as follows.

See the following table for information.

*Table 130. Settings on the local Windows system for a remote queue manager platform*

| Remote platform              | TPNAME                                                                                                                            | TPPATH                                    |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| z/OS or MVS/ESA without CICS | The same as in the corresponding side information about the remote queue manager.                                                 | -                                         |
| z/OS or MVS/ESA using CICS   | CKRC (sender) CKSV (requester) CKRC (server)                                                                                      | -                                         |
| IBM i                        | The same as the compare value in the routing entry on the IBM i system.                                                           | -                                         |
| UNIX and Linux systems       | The same as in the corresponding side information about the remote queue manager.                                                 | <i>MQ_INSTALLATION_PATH</i> /bin/amqcrs6a |
| Windows                      | As specified in the Windows Run Listener command, or the invocable Transaction Program that was defined using TpSetup on Windows. | <i>MQ_INSTALLATION_PATH</i> \bin\amqcrs6a |

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

If you have more than one queue manager on the same machine, ensure that the TPnames in the channel definitions are unique.

For the latest information about configuring AnyNet<sup>®</sup> SNA over TCP/IP, see the following online IBM documentation: AnyNet SNA over TCP/IP and SNA Node Operations.

### Related concepts:

“Sending end on LU 6.2”

Create a CPI-C side object (symbolic destination) from the administration application of the LU 6.2 product you are using. Enter this name in the Connection name field in the channel definition. Also create an LU 6.2 link to the partner.

“Receiving on LU 6.2”

Receiving channel programs are started in response to a startup request from the sending channel.

*Sending end on LU 6.2:*

Create a CPI-C side object (symbolic destination) from the administration application of the LU 6.2 product you are using. Enter this name in the Connection name field in the channel definition. Also create an LU 6.2 link to the partner.

In the CPI-C side object enter the partner LU Name at the receiving machine, the TP Name and the Mode Name. For example:

|                 |         |
|-----------------|---------|
| Partner LU Name | OS2R0G2 |
| Partner TP Name | recv    |
| Mode Name       | #INTER  |

*Receiving on LU 6.2:*

Receiving channel programs are started in response to a startup request from the sending channel.

To start a receiving channel program, a listener program has to be started to detect incoming network requests and start the associated channel. You start this listener program with the RUNMQLSR command, giving the TpName to listen on. Alternatively, you can use TpStart under SNA Server for Windows.

### Using the RUNMQLSR command

Example of the command to start the listener:

```
RUNMQLSR -t LU62 -n RECV [-m QMNAME]
```

where RECV is the TpName that is specified at the other (sending) end as the "TpName to start on the remote side". The last part in square brackets is optional and is not required for the default queue manager.

It is possible to have more than one queue manager running on one machine. You must assign a different TpName to each queue manager, and then start a listener program for each one. For example:

```
RUNMQLSR -t LU62 -m QM1 -n TpName1
RUNMQLSR -t LU62 -m QM2 -n TpName2
```

For the best performance, run the IBM MQ listener as a trusted application as described in Running channels and listeners as trusted applications. See Restrictions for trusted applications for information about trusted applications.

You can stop all IBM MQ listeners running on a queue manager that is inactive, using the command:

```
ENDMQLSR [-m QMNAME]
```

### Using Microsoft SNA Server on Windows

You can use TpSetup (from the SNA Server SDK) to define an invocable TP that then drives amqcrs6a.exe, or you can set various registry values manually. The parameters that should be passed to amqcrs6a.exe are:

```
-m QM -n TpName
```

where *QM* is the Queue Manager name and *TpName* is the TP Name. See the *Microsoft SNA Server APPC Programmers Guide* or the *Microsoft SNA Server CPI-C Programmers Guide* for more information.

If you do not specify a queue manager name, the default queue manager is assumed.

### Defining a NetBIOS connection on Windows:

IBM MQ uses three types of NetBIOS resource when establishing a NetBIOS connection to another IBM MQ product: sessions, commands, and names. Each of these resources has a limit, which is established either by default or by choice during the installation of NetBIOS.

Each running channel, regardless of type, uses one NetBIOS session and one NetBIOS command. The IBM NetBIOS implementation allows multiple processes to use the same local NetBIOS name. Therefore, only one NetBIOS name needs to be available for use by IBM MQ. Other vendors' implementations, for example Novell's NetBIOS emulation, require a different local name per process. Verify your requirements from the documentation for the NetBIOS product you are using.

In all cases, ensure that sufficient resources of each type are already available, or increase the maximums specified in the configuration. Any changes to the values require a system restart.

During system startup, the NetBIOS device driver displays the number of sessions, commands, and names available for use by applications. These resources are available to any NetBIOS-based application that is running on the same system. Therefore, it is possible for other applications to consume these resources before IBM MQ needs to acquire them. Your LAN network administrator should be able to clarify this for you.

#### Related concepts:

“Defining the IBM MQ local NetBIOS name”

The local NetBIOS name used by IBM MQ channel processes can be specified in three ways.

“Establishing the queue manager NetBIOS session, command, and name limits” on page 871

The queue manager limits for NetBIOS sessions, commands, and names can be specified in two ways.

“Establishing the LAN adapter number” on page 871

For channels to work successfully across NetBIOS, the adapter support at each end must be compatible. IBM MQ allows you to control the choice of LAN adapter (LANA) number by using the AdapterNum value in the NETBIOS stanza of your qm.ini file and by specifying the -a parameter on the runmqslr command.

“Initiating the NetBIOS connection” on page 872

Defining the steps needed to initiate a connection.

“Target listener for the NetBIOS connection” on page 872

Defining the steps to be undertaken at the receiving end of the NetBIOS connection.

#### *Defining the IBM MQ local NetBIOS name:*

The local NetBIOS name used by IBM MQ channel processes can be specified in three ways.

In order of precedence the three ways are:

1. The value specified in the -l parameter of the RUNMQLSR command, for example:

```
RUNMQLSR -t NETBIOS -l my_station
```

2. The MQNAME environment variable with a value that is established by the command:

```
SET MQNAME= my_station
```

You can set the MQNAME value for each process. Alternatively, you can set it at a system level in the Windows registry.

If you are using a NetBIOS implementation that requires unique names, you must issue a SET MQNAME command in each window in which an IBM MQ process is started. The MQNAME value is arbitrary but it must be unique for each process.

3. The NETBIOS stanza in the queue manager configuration file qm.ini. For example:

```
NETBIOS:
```

```
LocalName= my_station
```

**Note:**

1. Due to the variations in implementation of the NetBIOS products supported, you are advised to make each NetBIOS name unique in the network. If you do not, unpredictable results might occur. If you have problems establishing a NetBIOS channel and there are error messages in the queue-manager error log showing a NetBIOS return code of X'15', review your use of NetBIOS names.
2. On Windows, you cannot use your machine name as the NetBIOS name because Windows already uses it.
3. Sender channel initiation requires that a NetBIOS name be specified either by using the MQNAME environment variable or the LocalName in the qm.ini file.

*Establishing the queue manager NetBIOS session, command, and name limits:*

The queue manager limits for NetBIOS sessions, commands, and names can be specified in two ways.

In order of precedence these ways are:

1. The values specified in the RUNMQLSR command:

```
-s Sessions
-e Names
-o Commands
```

If the -m operand is not specified in the command, the values apply only to the default queue manager.

2. The NETBIOS stanza in the queue manager configuration file qm.ini. For example:

```
NETBIOS:
```

```
NumSess= Qmgr_max_sess
NumCmds= Qmgr_max_cmds
NumNames= Qmgr_max_names
```

*Establishing the LAN adapter number:*

For channels to work successfully across NetBIOS, the adapter support at each end must be compatible. IBM MQ allows you to control the choice of LAN adapter (LANA) number by using the AdapterNum value in the NETBIOS stanza of your qm.ini file and by specifying the -a parameter on the runmqsr command.

The default LAN adapter number used by IBM MQ for NetBIOS connections is 0. Verify the number being used on your system as follows:

On Windows, it is not possible to query the LAN adapter number directly through the operating system. Instead, you use the LANACFG.EXE command-line utility, available from Microsoft. The output of the tool shows the virtual LAN adapter numbers and their effective bindings. For further information about LAN adapter numbers, see the Microsoft Knowledge Base article 138037 *HOWTO: Use LANA Numbers in a 32-bit Environment*.

Specify the correct value in the NETBIOS stanza of the queue manager configuration file, qm.ini:

```
NETBIOS:
AdapterNum= n
```

where n is the correct LAN adapter number for this system.

*Initiating the NetBIOS connection:*

Defining the steps needed to initiate a connection.

To initiate the connection, follow these steps at the sending end:

1. Define the NetBIOS station name using the MQNAME or LocalName value.
2. Verify the LAN adapter number being used on your system and specify the correct file using the AdapterNum.
3. In the ConnectionName field of the channel definition, specify the NetBIOS name being used by the target listener program. On Windows, NetBIOS channels *must* be run as threads. Do this by specifying MCATYPE(THREAD) in the channel definition.

```
DEFINE CHANNEL (cname) CHLTYPE(SDR) +
TRPTYPE(NETBIOS) +
CONNNAME(your_station) +
XMITQ(xmitq) +
MCATYPE(THREAD) +
REPLACE
```

*Target listener for the NetBIOS connection:*

Defining the steps to be undertaken at the receiving end of the NetBIOS connection.

At the receiving end, follow these steps:

1. Define the NetBIOS station name using the MQNAME or LocalName value.
2. Verify the LAN adapter number being used on your system and specify the correct file using the AdapterNum.
3. Define the receiver channel:

```
DEFINE CHANNEL (cname) CHLTYPE(RCVR) +
TRPTYPE(NETBIOS) +
REPLACE
```

4. Start the IBM MQ listener program to establish the station and make it possible to contact it. For example:

```
RUNMQLSR -t NETBIOS -l your_station [-m qmgr]
```

This command establishes your\_station as a NetBIOS station waiting to be contacted. The NetBIOS station name must be unique throughout your NetBIOS network.

For the best performance, run the IBM MQ listener as a trusted application as described in “Running channels and listeners as trusted applications” on page 855. See Restrictions for trusted applications for information about trusted applications.

You can stop all IBM MQ listeners running on a queue manager that is inactive, using the command:

```
ENDMQLSR [-m QMNAME]
```

If you do not specify a queue manager name, the default queue manager is assumed.



## Setting up communication on UNIX and Linux systems

DQM is a remote queuing facility for IBM MQ. It provides channel control programs for the queue manager which form the interface to communication links, controllable by the system operator. The channel definitions held by distributed-queuing management use these connections.

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. To succeed, it is necessary for the connection to be defined and available. This section explains how to do this. You might also find it helpful to refer to the following sections:

- Example configuration - IBM MQ for AIX
- Example configuration - IBM MQ for HP-UX
- Example configuration - IBM MQ for Solaris
- Example configuration - IBM MQ for Linux

For Windows, see “Setting up communication for Windows” on page 865.

You can choose between two forms of communication for IBM MQ on UNIX and Linux systems:

- “Defining a TCP connection on UNIX and Linux”
- “Defining an LU 6.2 connection on UNIX and Linux” on page 877

Each channel definition must specify one only as the transmission protocol (Transport Type) attribute. One or more protocols can be used by a queue manager.

For IBM MQ MQI clients, it might be useful to have alternative channels using different transmission protocols. For more information about IBM MQ MQI clients, see Overview of IBM MQ MQI clients.

### Related concepts:

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Monitoring and controlling channels on Windows, UNIX and Linux platforms” on page 856

For DQM you need to create, monitor, and control the channels to remote queue managers. You can control channels using commands, programs, IBM MQ Explorer, files for the channel definitions, and a storage area for synchronization information.

“Configuring connections between the server and client” on page 696

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

### Defining a TCP connection on UNIX and Linux:

The channel definition at the sending end specifies the address of the target. The listener or inet daemon is configured for the connection at the receiving end.

### Sending end

Specify the host name, or the TCP address of the target machine, in the Connection Name field of the channel definition. The port to connect to defaults to 1414. Port number 1414 is assigned by the Internet Assigned Numbers Authority to IBM MQ.

To use a port number other than the default, change the connection name field thus:

```
Connection Name REMHOST(1822)
```

where REMHOST is the host name of the remote machine and 1822 is the port number required. (This must be the port that the listener at the receiving end is listening on.)

Alternatively you can change the port number by specifying it in the queue manager configuration file (qm.ini):

```
TCP:
Port=1822
```

For more information about the values you set using qm.ini, see Configuration file stanzas for distributed queuing.

## Receiving on TCP

You can use either the TCP/IP listener, which is the inet daemon (inetd), or the IBM MQ listener.

Some Linux distributions now use the extended inet daemon (xinetd) instead of the inet daemon. For information about how to use the extended inet daemon on a Linux system, see Establishing a TCP connection on Linux .

### Related concepts:

“Using the TCP/IP listener”

To start channels on UNIX and Linux, the /etc/services file and the inetd.conf file must be edited

“Using the TCP listener backlog option” on page 875

In TCP, connections are treated incomplete unless three-way handshake takes place between the server and the client. These connections are called outstanding connection requests. A maximum value is set for these outstanding connection requests and can be considered a backlog of requests waiting on the TCP port for the listener to accept the request.

“Using the IBM MQ listener” on page 876

To run the listener supplied with IBM MQ, which starts new channels as threads, use the runmq1sr command.

“Using the TCP/IP SO\_KEEPALIVE option” on page 877

On some UNIX and Linux systems, you can define how long TCP waits before checking that the connection is still available, and how frequently it tries the connection again if the first check fails. This is either a kernel tunable parameter, or can be entered at the command line.

*Using the TCP/IP listener:*

To start channels on UNIX and Linux, the /etc/services file and the inetd.conf file must be edited

Follow these instructions:

1. Edit the /etc/services file:

**Note:** To edit the /etc/services file, you must be logged in as a superuser or root. You can change this, but it must match the port number specified at the sending end.

Add the following line to the file:

```
MQSeries 1414/tcp
```

where 1414 is the port number required by IBM MQ. The port number must not exceed 65535.

2. Add a line in the inetd.conf file to call the program amqcrsta, where `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed:

```
MQSeries stream tcp nowait mqm MQ_INSTALLATION_PATH/bin/amqcrsta amqcrsta
[-m Queue_Man_Name]
```

The updates are active after inetd has reread the configuration files. To do this, issue the following commands from the root user ID:

- On AIX:

```
refresh -s inetd
```

- On HP-UX, from the mqm user ID:  
inetd -c
- On Solaris 10 or later:  
inetconv
- On other UNIX and Linux systems (including Solaris 9):  
kill -1 < process number >

When the listener program started by inetd inherits the locale from inetd, it is possible that the MQMDE is not honored (merged) and is placed on the queue as message data. To ensure that the MQMDE is honored, you must set the locale correctly. The locale set by inetd might not match that chosen for other locales used by IBM MQ processes. To set the locale:

1. Create a shell script which sets the locale environment variables LANG, LC\_COLLATE, LC\_CTYPE, LC\_MONETARY, LC\_NUMERIC, LC\_TIME, and LC\_MESSAGES to the locale used for other IBM MQ process.
2. In the same shell script, call the listener program.
3. Modify the inetd.conf file to call your shell script in place of the listener program.

It is possible to have more than one queue manager on the server. You must add a line to each of the two files, for each of the queue managers. For example:

```
MQSeries1 1414/tcp
MQSeries2 1822/tcp
MQSeries2 stream tcp nowait mqm MQ_INSTALLATION_PATH/bin/amqcrsta amqcrsta -m QM2
```

Where *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

This avoids error messages being generated if there is a limitation on the number of outstanding connection requests queued at a single TCP port. For information about the number of outstanding connection requests, see “Using the TCP listener backlog option.”

*Using the TCP listener backlog option:*

In TCP, connections are treated incomplete unless three-way handshake takes place between the server and the client. These connections are called outstanding connection requests. A maximum value is set for these outstanding connection requests and can be considered a backlog of requests waiting on the TCP port for the listener to accept the request.

The default listener backlog values are shown in Table 131.

*Table 131. Maximum outstanding connection requests queued at a TCP/IP port*

| Server platform     | Maximum connection requests |
|---------------------|-----------------------------|
| AIX                 | 100                         |
| HP-UX               | 20                          |
| Linux               | 100                         |
| IBM i               | 255                         |
| Solaris             | 100                         |
| Windows Server      | 100                         |
| Windows Workstation | 100                         |
| z/OS                | 255                         |

If the backlog reaches the values shown in Table 131, the TCP/IP connection is rejected and the channel is not able to start.

For MCA channels, this results in the channel going into a RETRY state and trying the connection again at a later time.

However, to avoid this error, you can add an entry in the qm.ini file:

```
TCP:
ListenerBacklog = n
```

This overrides the default maximum number of outstanding requests (see Table 131 on page 875 ) for the TCP/IP listener.

**Note:** Some operating systems support a larger value than the default. If necessary, this value can be used to avoid reaching the connection limit.

To run the listener with the backlog option switched on either:

- Use the `runmq1sr -b` command, or
- Use the MQSC command **DEFINE LISTENER** with the BACKLOG attribute set to the required value.

For information about the `runmq1sr` command, see `runmq1sr`. For information about the `DEFINE LISTENER` command, see the `DEFINE LISTENER`.

*Using the IBM MQ listener:*

To run the listener supplied with IBM MQ, which starts new channels as threads, use the `runmq1sr` command.

For example:

```
runmq1sr -t tcp [-m QMNAME] [-p 1822]
```

The square brackets indicate optional parameters; QMNAME is not required for the default queue manager, and the port number is not required if you are using the default (1414). The port number must not exceed 65535.

For the best performance, run the IBM MQ listener as a trusted application as described in “Running channels and listeners as trusted applications” on page 855. See Restrictions for trusted applications for information about trusted applications.

You can stop all IBM MQ listeners running on a queue manager that is inactive, using the command:

```
endmq1sr [-m QMNAME]
```

If you do not specify a queue manager name, the default queue manager is assumed.

Using the TCP/IP SO\_KEEPALIVE option:

On some UNIX and Linux systems, you can define how long TCP waits before checking that the connection is still available, and how frequently it tries the connection again if the first check fails. This is either a kernel tunable parameter, or can be entered at the command line.

If you want to use the SO\_KEEPALIVE option (for more information, see “Checking that the other end of the channel is still available” on page 840 ) you must add the following entry to your queue manager configuration file (qm.ini):

```
TCP:
KeepAlive=yes
```

See the documentation for your UNIX and Linux system for more information.

### Defining an LU 6.2 connection on UNIX and Linux:

SNA must be configured so that an LU 6.2 conversation can be established between the two machines.

For the latest information about configuring SNA over TCP/IP, see the following online IBM documentation: Communications Server.

SNA must be configured so that an LU 6.2 conversation can be established between the two systems.

See the *Multiplatform APPC Configuration Guide* and the following table for information.

Table 132. Settings on the local UNIX and Linux system for a remote queue manager platform

| Remote platform        | TPNAME                                                                                                                            | TPPATH                                    |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| z/OS without CICS      | The same as the corresponding TPName in the side information about the remote queue manager.                                      | -                                         |
| z/OS using CICS        | CKRC (sender) CKSV (requester) CKRC (server)                                                                                      | -                                         |
| IBM i                  | The same as the compare value in the routing entry on the IBM i system.                                                           | -                                         |
| UNIX and Linux systems | The same as the corresponding TPName in the side information about the remote queue manager.                                      | <i>MQ_INSTALLATION_PATH</i> /bin/amqcrs6a |
| Windows                | As specified in the Windows Run Listener command, or the invocable Transaction Program that was defined using TpSetup on Windows. | <i>MQ_INSTALLATION_PATH</i> \bin\amqcrs6a |

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

If you have more than one queue manager on the same machine, ensure that the TPnames in the channel definitions are unique.

### Related concepts:

#### “Sending end”

On UNIX and Linux systems, create a CPI-C side object (symbolic destination) and enter this name in the Connection name field in the channel definition. Also create an LU 6.2 link to the partner.

#### “Receiving on LU 6.2”

On UNIX and Linux systems, create a listening attachment at the receiving end, an LU 6.2 logical connection profile, and a TPN profile.

#### *Sending end:*

On UNIX and Linux systems, create a CPI-C side object (symbolic destination) and enter this name in the Connection name field in the channel definition. Also create an LU 6.2 link to the partner.

In the CPI-C side object enter the partner LU name at the receiving machine, the transaction program name and the mode name. For example:

|                             |         |
|-----------------------------|---------|
| Partner LU Name             | REMHOST |
| Remote TP Name              | recv    |
| Service Transaction Program | no      |
| Mode Name                   | #INTER  |

On HP-UX, use the APPCLLU environment variable to name the local LU that the sender should use. On Solaris, set the APPC\_LOCAL\_LU environment variable to be the local LU name.

SECURITY PROGRAM is used, where supported by CPI-C, when IBM MQ attempts to establish an SNA session.

#### *Receiving on LU 6.2:*

On UNIX and Linux systems, create a listening attachment at the receiving end, an LU 6.2 logical connection profile, and a TPN profile.

In the TPN profile, enter the full path to the executable file and the Transaction Program name:

|                             |                                           |
|-----------------------------|-------------------------------------------|
| Full path to TPN executable | <i>MQ_INSTALLATION_PATH</i> /bin/amqcrs6a |
| Transaction Program name    | recv                                      |
| User ID                     | 0                                         |

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

On systems where you can set the user ID, specify a user who is a member of the mqm group. On AIX, Solaris, and HP-UX, set the APPCTPN (transaction name) and APPCLLU (local LU name) environment variables (you can use the configuration panels for the invoked transaction program).

You might need to use a queue manager other than the default queue manager. If so, define a command file that calls:

```
amqcrs6a -m Queue_Man_Name
```

then call the command file.

## Monitoring and controlling channels on IBM i

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers. Each queue manager has a DQM program for controlling interconnections to compatible remote queue managers.

The following list is a brief description of the components of the channel control function:

- Channel definitions are held as queue manager objects.
- The channel commands are a subset of the IBM MQ for IBM i set of commands.  
Use the command GO CMDMQM to display the full set of IBM MQ for IBM i commands.
- You use channel definition panels, or commands to:
  - Create, copy, display, change, and delete channel definitions
  - Start and stop channels, ping, reset channel sequence numbers, and resolve in-doubt messages when links cannot be re-established
  - Display status information about channels
- Channels can also be managed using MQSC
- Channels can also be managed using IBM MQ Explorer
- Sequence numbers and *logical unit of work (LUW)* identifiers are stored in the synchronization file, and are used for channel synchronization purposes.

You can use the commands and panels to: define message channels and associated objects, and monitor and control message channels. By using the F4=Prompt key, you can specify the relevant queue manager. If you do not use the prompt, the default queue manager is assumed. With F4=Prompt, an additional panel is displayed where you can enter the relevant queue manager name and sometimes other data.

The objects you need to define with the panels are:

- Transmission queues
- Remote queue definitions
- Queue manager alias definitions
- Reply-to queue alias definitions
- Reply-to local queues
- Message channel definitions

For more information about the concepts involved in the use of these objects, see “Configuring distributed queuing” on page 802.

Channels must be completely defined, and their associated objects must exist and be available for use, before a channel can be started.

In addition, the particular communication link for each channel must be defined and available before a channel can be run. For a description of how LU 6.2 and TCP/IP links are defined, see the particular communication guide for your installation.

For more information about creating and working with objects, see:

- “Creating objects” on page 880
- “Creating a channel” on page 880
- “Starting a channel” on page 883
- “Selecting a channel” on page 883
- “Browsing a channel” on page 884
- “Renaming a channel” on page 886
- “Work with channel status” on page 886

- “Work-with-channel choices” on page 887

**Related concepts:**

“Setting up communication for IBM MQ for IBM i” on page 894

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. For it to succeed, it is necessary for the connection to be defined and available.

“Configuring connections between the server and client” on page 696

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

**Related information:**

Example configuration - IBM MQ for IBM i

Message channel planning example for IBM MQ for IBM i

IBM MQ for IBM i CL commands

**Creating objects**

You can use the CRTMQMQ command to create the queue and alias objects.

You can create the queue and alias objects, such as: transmission queues, remote queue definitions, queue manager alias definitions, reply-to queue alias definitions, and reply-to local queues.

For a list of default objects, see IBM MQ for IBM i system and default objects.

**Creating a channel**

You can create a channel from the Create Channel panel or by using the CRTMQMCHL command on the command line.

To create a channel:

1. Use F6 from the Work with MQM Channels panel (WRKMQMCHL).  
Alternatively, use the CRTMQMCHL command from the command line.  
Either way, the Create Channel panel is displayed. Type:
  - The name of the channel in the field provided
  - The channel type for this end of the link
2. Press enter.

**Note:** You must name all the channels in your network uniquely. As shown in Network diagram showing all channels, including the source and target queue manager names in the channel name is a good way to do so.

Your entries are validated and errors are reported immediately. Correct any errors and continue.

You are presented with the appropriate channel settings panel for the type of channel you have chosen. Complete the fields with the information you have gathered previously. Press enter to create the channel.

You are provided with help in deciding on the content of the various fields in the descriptions of the channel definition panels in the help panels, and in Channel attributes.



```

Create MQM Channel (CRTMQMCHL)

Type choices, press Enter.

Channel name > CHANNAME_____
Channel type > *SDR__ *RCVR, *SDR, *SVR, *RQSTR...
Message Queue Manager name *DFT_____

Replace *NO_ *NO, *YES
Transport type *TCP_____ *LU62, *TCP, *SYSDFTCHL
Text 'description' > 'Example Channel Definition'_____

Connection name *SYSDFTCHL_____

More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 105. Create channel (1)

```

Create MQM Channel (CRTMQMCHL)

Type choices, press Enter.

Transmission queue 'TRANSMISSION_QUEUE_NAME'_____

Message channel agent *NONE_____ Name, *SYSDFTCHL, *NONE
Library _____ Name
Message channel agent user ID . *SYSDFTCHL_ Character value...
Coded Character Set Identifier *SYSDFTCHL_ 0-9999, *SYSDFTCHL
Batch size 50_____ 1-9999, *SYSDFTCHL
Disconnect interval 6000_____ 1-999999, *SYSDFTCHL
Short retry interval 60_____ 0-999999999, *SYSDFTCHL
Short retry count 10_____ 0-999999999, *SYSDFTCHL
Long retry interval 1200_____ 0-999999999, *SYSDFTCHL
Long retry count 999999999_ 0-999999999, *SYSDFTCHL
Security exit *NONE_____ Name, *SYSDFTCHL, *NONE
Library _____ Name
Security exit user data *SYSDFTCHL_____

More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 106. Create channel (2)

Create MQM Channel (CRTMQMCHL)

Type choices, press Enter.

```
Send exit *NONE_____ Name, *SYSDFTCHL, *NONE
Library _____ Name
+ for more values _____
Send exit user data _____
+ for more values _____
Receive exit *NONE_____ Name, *SYSDFTCHL, *NONE
Library _____ Name
+ for more values _____
Receive exit user data _____
+ for more values _____
Message exit *NONE_____ Name, *SYSDFTCHL, *NONE
Library _____ Name
+ for more values _____
More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
```

Figure 107. Create channel (3)

Create MQM Channel (CRTMQMCHL)

Type choices, press Enter.

```
Message exit user data _____
+ for more values _____
Convert message *SYSDFTCHL_ *YES, *NO, *SYSDFTCHL
Sequence number wrap 999999999_ 100-999999999, *SYSDFTCHL
Maximum message length 4194304_ 0-4194304, *SYSDFTCHL
Heartbeat interval 300_ 0-999999999, *SYSDFTCHL
Non Persistent Message Speed . . *FAST_____ *FAST, *NORMAL, *SYSDFTCHL
Password *SYSDFTCHL_ Character value, *BLANK...
Task User Profile *SYSDFTCHL_ Character value, *BLANK...
Transaction Program Name *SYSDFTCHL

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
```

Figure 108. Create channel (4)

## Starting a channel

You can start a channel from the Work with Channels panel or by using the STRMQMCHL command on the command line.

Listeners are valid for TCP only. For SNA listeners, you must configure your communications subsystem.

For applications to be able to exchange messages, you must start a listener program for inbound connections using the STRMQMLSR command.

For outbound connections, you must start the channel in one of the following ways:

1. Use the CL command STRMQMCHL, specifying the channel name, to start the channel as a process or a thread, depending on the MCATYPE parameter. (If channels are started as threads, they are threads of a channel initiator.)  
STRMQMCHL CHLNAME(QM1.TO.QM2) MQNAME(MYQMGR)
2. Use a channel initiator to trigger the channel. One channel initiator is started automatically when the queue manager is started. This automatic start can be eliminated by changing the chinit stanza in the qm.ini file for that queue manager.
3. Use the WRKMQMCHL command to begin the Work with Channels panel and choose option 14 to start a channel.

## Selecting a channel

You can select a channel from the Work With channels panel.

To select a channel, use the WRKMQMCHL command to begin at the Work with Channels panel:

1. Move the cursor to the option field associated with the required channel name.
2. Type an option number.
3. Press enter to activate your choice.

If you select more than one channel, the options are activated in sequence.

```
Work with MQM Channels

Queue Manager Name . . : CNX

Type options, press Enter.
2=Change 3=Copy 4=Delete 5=Display 8=Work with Status 13=Ping
14=Start 15=End 16=Reset 17=Resolve

Opt Name Type Transport Status
CHLNIC *RCVR *TCP INACTIVE
CORSAIR.TO.MUSTANG *SDR *LU62 INACTIVE
FV.CHANNEL.MC.DJE1 *RCVR *TCP INACTIVE
FV.CHANNEL.MC.DJE2 *SDR *TCP INACTIVE
FV.CHANNEL.MC.DJE3 *RQSTR *TCP INACTIVE
FV.CHANNEL.MC.DJE4 *SVR *TCP INACTIVE
FV.CHANNEL.PETER *RCVR *TCP INACTIVE
FV.CHANNEL.PETER.LU *RCVR *LU62 INACTIVE
FV.CHANNEL.PETER.LU1 *RCVR *LU62 INACTIVE
More...
Parameters or command
====
F3=Exit F4=Prompt F5=Refresh F6=Create F9=Retrieve F12=Cancel
F21=Print
```

Figure 109. Work with channels

## Browsing a channel

You can browse a channel from the Display Channel panel or by using the DSPMQMCHL command on the command line.

To browse the settings of a channel, use the WRKMQMCHL command to begin at the Display Channel panel:

1. Type option 5 (Display) against the required channel name.
2. Press enter to activate your choice.

If you select more than one channel, they are presented in sequence.

Alternatively, you can use the DSPMQMCHL command from the command line.

This results in the appropriate Display Channel panel being displayed with details of the current settings for the channel. The fields are described in Channel attributes.

```
Display MQM Channel

Channel name : ST.JST.2T01
Queue Manager Name : QMREL
Channel type : *SDR
Transport type : *TCP
Text 'description' : John's sender to WINSDOA1

Connection name : MUSTANG

Transmission queue : WINSDOA1

Message channel agent :
Library :
Message channel agent user ID : *NONE
Batch interval : 0
Batch size : 50
Disconnect interval : 6000

F3=Exit F12=Cancel F21=Print
```

Figure 110. Display a TCP/IP channel (1)

```

Display MQM Channel

Short retry interval : 60
Short retry count : 10
Long retry interval : 6000
Long retry count : 10
Security exit :
Library :
Security exit user data . . . :
Send exit :
Library :
Send exit user data :
Receive exit :
Library :
Receive exit user data :
Message exit :
Library :
Message exit user data :
More...

F3=Exit F12=Cancel F21=Print

```

Figure 111. Display a TCP/IP channel (2)

```

Display MQM Channel

Sequence number wrap : 999999999
Maximum message length : 10000
Convert message : *NO
Heartbeat interval : 300
Nonpersistent message speed . . *FAST

Bottom

F3=Exit F12=Cancel F21=Print

```

Figure 112. Display a TCP/IP channel (3)

## Renaming a channel

You can rename a channel from the Work with Channels panel.

To rename a message channel, begin at the Work with Channels panel:

1. End the channel.
2. Use option 3 (Copy) to create a duplicate with the new name.
3. Use option 5 (Display) to check that it has been created correctly.
4. Use option 4 (Delete) to delete the original channel.

If you decide to rename a message channel, ensure that both channel ends are renamed at the same time.

## Work with channel status

You can work with the channel status from the Work with Channel Status panel.

Use the WRKMQMCHST command to display the first of a set of panels showing the status of your channels. You can view the status panels in sequence when you select Change-view (F11).

Alternatively, selecting option 8 (Work with Status) from the Work with MQM Channels panel also displays the first status panel.

```
MQSeries Work with Channel Status

Type options, press Enter.
5=Display 13=Ping 14=Start 15=End 16=Reset 17=Resolve

Opt Name Connection Indoubt Last Seq
CARTS_CORSAIR_CHAN GBIBMIYA.WINSDOA1 NO 1
CHLNIC 9.20.2.213 NO 3
FV.CHANNEL.PETER2 9.20.2.213 NO 6225
JST.1.1.2 9.20.2.201 NO 28
MP_MUST_TO_CORS 9.20.2.213 NO 100
MUSTANG.TO.CORSAIR GBIBMIYA.WINSDOA1 NO 10
MP_CORS_TO_MUST 9.20.2.213 NO 101
JST.2.3 9.5.7.126 NO 32
PF_WINSDOA1_LU62 GBIBMIYA.IYA80020 NO 54
PF_WINSDOA1_LU62 GBIBMIYA.WINSDOA1 NO 500
ST.JCW.EXIT.2T01.CHL 9.20.2.213 NO 216

Bottom
Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Create F9=Retrieve F11=Change view
F12=Cancel F21=Print
```

Figure 113. First of the set of channel status panels

The options available in the Work with Channel Status panel are:

| Menu option | Description                                       |
|-------------|---------------------------------------------------|
| 5=Display   | Displays the channel settings.                    |
| 13=Ping     | Initiates a Ping action, where appropriate.       |
| 14=Start    | Starts the channel.                               |
| 15=End      | Stops the channel.                                |
| 16=Reset    | Resets the channel sequence number.               |
| 17=Resolve  | Resolves an in-doubt channel situation, manually. |

## Work-with-channel choices

The Work with Channels panel is reached with the command `WRKMQMCHL`, and it allows you to monitor the status of all channels listed, and to issue commands against selected channels.

The options available in the Work with Channel panel are:

| Menu option                        | Description                                                                                                                                  |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| "2=Change"                         | Changes the attributes of a channel.                                                                                                         |
| "3=Copy" on page 888               | Copies the attributes of a channel to a new channel.                                                                                         |
| "4=Delete" on page 888             | Deletes a channel.                                                                                                                           |
| "5=Display" on page 888            | Displays the current settings for the channel.                                                                                               |
| "6=Create" on page 888             | Displays the Create channel panel                                                                                                            |
| "8=Work with Status" on page 889   | Displays the channel status panels.                                                                                                          |
| "13=Ping" on page 890              | Runs the Ping facility to test the connection to the adjacent system by exchanging a fixed data message with the remote end.                 |
| "14=Start" on page 890             | Starts the selected channel, or resets a disabled receiver channel.                                                                          |
| "15=End" on page 891               | Requests the channel to close down.                                                                                                          |
| "16=Reset" on page 892             | Requests the channel to reset the sequence numbers on this end of the link. The numbers must be equal at both ends for the channel to start. |
| "17=Resolve" on page 892           | Requests the channel to resolve in-doubt messages without establishing connection to the other end.                                          |
| "18=Display authority" on page 893 | Displays IBM MQ object authority                                                                                                             |
| "19=Grant authority" on page 893   | Grants IBM MQ object authority                                                                                                               |
| "20=Revoke authority" on page 893  | Revokes IBM MQ object authority                                                                                                              |
| "21=Recover object" on page 893    | Recovers IBM MQ object                                                                                                                       |
| "22=Record image" on page 893      | Records IBM MQ object image                                                                                                                  |

### 2=Change:

Use the Change option to change an existing channel definition.

The Change option, or the `CHGMQMCHL` command, changes an existing channel definition, except for the channel name. Type over the fields to be changed in the channel definition panel, and then save the updated definition by pressing enter.

### **3=Copy:**

Use the Copy option to copy an existing channel.

The Copy option uses the CPYMQMCHL command to copy an existing channel. The Copy panel enables you to define the new channel name. However, you must restrict the characters used to those characters that are valid for IBM MQ for IBM i object names; see the *IBM MQ for IBM i System Administration*.

Press enter on the Copy panel to display the details of current settings. You can change any of the new channel settings. Save the new channel definition by pressing enter.

### **4=Delete:**

Use the Delete option to delete the selected channel.

A panel is displayed to confirm or cancel your request.

### **5=Display:**

Use the Display option to display the current definitions for the channel.

This choice displays the panel with the fields showing the current values of the parameters, and protected against user input.

### **6=Create:**

Use the Create option to display the Create channel panel.

Use the Create option, or enter the CRTMQMCHL command from the command line, to obtain the Create Channel panel. There are examples of Create Channel panels, starting at Figure 105 on page 881.

With this panel, you create a channel definition from a screen of fields filled with default values supplied by IBM MQ for IBM i. Type the name of the channel, select the type of channel you are creating, and the communication method to be used.

When you press enter, the panel is displayed. Type information in all the required fields in this panel, and the remaining panels, and then save the definition by pressing enter.

The channel name must be the same at both ends of the channel, and unique within the network. However, you must restrict the characters used to those characters that are valid for IBM MQ for IBM i object names.

All panels have default values supplied by IBM MQ for IBM i for some fields. You can customize these values, or you can change them when you are creating or copying channels. To customize the values, see the *IBM MQ for IBM i System Administration*.

You can create your own set of channel default values by setting up dummy channels with the required defaults for each channel type, and copying them each time you want to create new channel definitions.



**Related information:**

Channel attributes

**8=Work with Status:**

Use Work with Status to see detailed channel status information.

The status column tells you whether the channel is active or inactive, and is displayed continuously in the Work with MQM Channels panel. Use option 8 (Work with Status) to see more status information displayed. Alternatively, this information can be displayed from the command line with the WRKMQMCHST command. See “Work with channel status” on page 886.

- Channel name
- Channel type
- Channel status
- Channel instance
- Remote queue manager
- Transmission queue name
- Communication connection name
- In-doubt status of channel
- Last sequence number
- Number of indoubt messages
- In-doubt sequence number
- Number of messages on transmission queue
- Logical unit of work identifier
- In-doubt logical unit of work identifier
- Channel substate
- Channel monitoring
- Header compression
- Message compression
- Compression time indicator
- Compression rate indicator
- Transmission queue time indicator
- Network time indicator
- Exit time indicator
- Batch size indicator
- Current shared conversations
- Maximum shared conversations

### **13=Ping:**

Use the Ping option to exchange a fixed data message with the remote end.

A successful IBM MQ Ping gives some confidence to the system supervisor that the channel is available and functioning.

Ping does not involve the use of transmission queues and target queues. It uses channel definitions, the related communication link, and the network setup.

It is available from sender and server channels, only. The corresponding channel is started at the far side of the link, and performs the start-up parameter negotiation. Errors are notified normally.

The result of the message exchange is presented in the Ping panel for you, and is the returned message text, together with the time the message was sent, and the time the reply was received.

### **Ping with LU 6.2**

When Ping is invoked in IBM MQ for IBM i, it is run with the user ID of the user requesting the function, whereas the normal way that a channel program is run is for the QMQM user ID to be taken for channel programs. The user ID flows to the receiving side and it must be valid on the receiving end for the LU 6.2 conversation to be allocated.

### **14=Start:**

Use the Start option to start a channel manually.

The Start option is available for sender, server, and requester channels. It is not necessary where a channel has been set up with queue manager triggering.

The Start option is also used for receiver, server-connection, cluster sender, and cluster receiver channels. Starting a receiver channel that is in STOPPED state means that it can be started from the remote channel.

When started, the sending MCA reads the channel definition file and opens the transmission queue. A channel start-up sequence is issued, which remotely starts the corresponding MCA of the receiver or server channel. When they have been started, the sender and server processes await messages arriving on the transmission queue and transmit them as they arrive.

When you use triggering, you must start the continuously running trigger process to monitor the initiation queue. The STRMQMCHLI command can be used for starting the process.

At the far end of a channel, the receiving process might be started in response to a channel startup from the sending end. The method of doing so is different for LU 6.2 and TCP/IP connected channels:

- LU 6.2 connected channels do not require any explicit action at the receiving end of a channel.
- TCP connected channels require a listener process to be running continuously. This process awaits channel startup requests from the remote end of the link and starts the process defined in the channel definitions for that connection.

When the remote system is IBM i, you can use the STRMQMLSR command.

Use of the Start option always causes the channel to resynchronize, where necessary.

For the start to succeed:

- Channel definitions, local and remote must exist. If there is no appropriate channel definition for a receiver or server-connection channel, a default one is created automatically if the channel is auto-defined. See Channel auto-definition exit program.

- The transmission queue must exist, be enabled for GETs, and have no other channels using it.
- MCAs, local and remote, must exist.
- The communication link must be available.
- The queue managers must be running, local and remote.
- The message channel must be inactive.

To transfer messages, remote queues and remote queue definitions must exist.

A message is returned to the panel confirming that the request to start a channel has been accepted. For confirmation that the Start process has succeeded, check the system log, or press F5 (refresh the screen).

### **15=End:**

Use End to stop channel activity

Use the End option to request the channel to stop activity. The channel does not send any more messages.

Select F4 before pressing enter to choose whether the channel becomes STOPPED or INACTIVE, and whether to stop the channel using a CONTROLLED or an IMMEDIATE stop. A stopped channel must be restarted by the operator to become active again. An inactive channel can be triggered.

### **Stop immediate**

Use Stop immediate to stop a channel without completing any unit of work.

This option terminates the channel process. As a result the channel does not complete processing the current batch of messages, and cannot, therefore, leave the channel in doubt. In general, it is better for the operators to use the controlled stop option.

### **Stop controlled**

Use Stop controlled to stop a channel at the end of the current unit of work.

This choice requests the channel to close down in an orderly way; the current batch of messages is completed, and the sync point procedure is carried out with the other end of the channel.

### **Restarting stopped channels**

When a channel goes into STOPPED state, you must restart the channel manually.

To do restart the channel, issue one of the following commands:

- The START CHANNEL MQSC command
- The Start Channel PCF command
- the IBM MQ Explorer
- other platform-specific mechanisms, as follows:

#### **For z/OS:**

The Start a channel panel

#### **For IBM i:**

The STRMQMCHL CL command or the START option on the WRKMQMCHL panel

For sender or server channels, when the channel entered the STOPPED state, the associated transmission queue was set to GET(DISABLED) and triggering was set off. When the start request is received, these attributes are reset automatically.

If the channel initiator (on z/OS ) or queue manager (on distributed platforms) stops while a channel is in RETRYING or STOPPED status, the channel status is remembered when the channel initiator or queue manager is restarted. However, the channel status for the SVRCONN channel type is reset if the channel initiator or queue manager stops while the channel is in STOPPED status.

#### **16=Reset:**

Use the Reset option to force a new message sequence.

The Reset option changes the message sequence number. Use it with care, and only after you have used the Resolve option to resolve any in-doubt situations. This option is available only at the sender or server channel. The first message starts the new sequence the next time the channel is started.

#### **17=Resolve:**

Use the Resolve option to force a local commit or backout of in-doubt messages held in a transmission queue.

Use the Resolve option when messages are held in-doubt by a sender or server, for example because one end of the link has terminated, and there is no prospect of it recovering. The Resolve option accepts one of two parameters: BACKOUT or COMMIT. Backout restores messages to the transmission queue, while Commit discards them.

The channel program does not try to establish a session with a partner. Instead, it determines the logical unit of work identifier (LUWID) which represents the in-doubt messages. It then issues, as requested, either:

- BACKOUT to restore the messages to the transmission queue; or
- COMMIT to delete the messages from the transmission queue.

For the resolution to succeed:

- The channel must be inactive
- The channel must be in doubt
- The channel type must be sender or server
- The channel definition, local, must exist
- The queue manager must be running, local

### **18=Display authority:**

Use the Display authority option to display what actions a user is authorized to perform on a specific IBM MQ object.

For a chosen object, and user, the DSPMQAUT command shows the authorizations the user has to perform actions on an IBM MQ object. If the user is a member of multiple groups, then the command shows the combined authorization of all the groups to the object.

### **19=Grant authority:**

Use the Grant authority option to grant the authority to perform actions on IBM MQ objects to another user or group of users.

The GRMQMAUT command is only available to users in the QMQMADM group. A user in QMQMADM grants authority to other users to perform actions on the IBM MQ objects named in the command either by identifying the users by name, or by granting authority to all users in \*PUBLIC.

### **20=Revoke authority:**

Use Revoke authority to remove authorization to perform actions on objects from users.

The RVKMQMAUT command is only available to users in the QMQMADM group. A user in the QMQMADM group removes the authority from other users to perform actions on the IBM MQ objects named in the command either by identifying the users by name, or by revoking authority from all users in \*PUBLIC.

### **21=Recover object:**

Use Recover object to restore damaged objects from information stored in IBM MQ journals.

Recover object uses the Re-create MQ Object command (RCRMQMOBJ) to recover all objects that are damaged named in the command. If an object is not damaged, then no action is performed on that object.

### **22=Record image:**

Use Record image to reduce the number of journal receivers required for the recovery of a set of objects, and to minimize recovery time.

The RCDMQMIMG command takes a checkpoint for all the objects that are selected in the command. It synchronizes the current values of the objects in the integrated file system (IFS) with later information about the objects, such as MQPUTs and MQGETs recorded in journal receivers.

When the command completes the objects in the IFS are up to date, and those journal receivers are no longer required to be present to recover the objects. Any disconnected journal receivers can be detached (as long as they are not required to be present to recover other objects).

## Setting up communication for IBM MQ for IBM i

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. For it to succeed, it is necessary for the connection to be defined and available.

DQM is a remote queuing facility for IBM MQ for IBM i. It provides channel control programs for the IBM MQ for IBM i queue manager which form the interface to communication links, controllable by the system operator.

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. For it to succeed, it is necessary for the connection to be defined and available. This section explains how to ensure that the connection is defined and available.

Before a channel can be started, the transmission queue must be defined as described in this section, and must be included in the message channel definition.

You can choose between the following two forms of communication between IBM MQ for IBM i systems:

- “Defining a TCP connection on IBM i”

For TCP, a host address can be used, and these connections are set up as described in the *IBM i Communication Configuration Reference*.

In the TCP environment, each distributed service is allocated a unique TCP address which can be used by remote machines to access the service. The TCP address consists of a host name/number and a port number. All queue managers use such a number to communicate with each other by way of TCP.

- “Receiving on TCP” on page 895

This form of communication requires the definition of an IBM i SNA logical unit type 6.2 (LU 6.2) that provides the physical link between the IBM i system serving the local queue manager and the system serving the remote queue manager. Refer to the *IBM i Communication Configuration Reference* for details on configuring communications in IBM i.

In addition, where needed, the triggering arrangement must be prepared with the definition of the necessary processes and queues.

### Related concepts:

“Monitoring and controlling channels on IBM i” on page 879

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers. Each queue manager has a DQM program for controlling interconnections to compatible remote queue managers.

### Related information:

Example configuration - IBM MQ for IBM i

Message channel planning example for IBM MQ for IBM i

Intercommunication jobs on IBM i

Channel states on IBM i

### Defining a TCP connection on IBM i:

You can define a TCP connection within the channel definition using the Connection Name field.

The channel definition contains a field, CONNECTION NAME, that contains either the TCP network address of the target or the host name (for example ABCHOST). The TCP network address can be in IPv4 dotted decimal form (for example 127.0.0.1) or IPv6 hexadecimal form (for example 2001:DB8:0:0:0:0:0). If the CONNECTION NAME is a host name or a name server, the IBM i host table is used to convert the host name into a TCP host address.

A port number is required for a complete TCP address; if this number is not supplied, the default port number 1414 is used. On the initiating end of a connection (sender, requester, and server channel types) it is possible to provide an optional port number for the connection, for example:

```
Connection name 127.0.0.1 (1555)
```

In this case the initiating end attempts to connect to a receiving program at port 1555.

#### **Related concepts:**

“Receiving on TCP”

Receiving channel programs are started in response to a startup request from the sending channel. To respond to the startup request, a listener program has to be started to detect incoming network requests and start the associated channel. You start this listener program with the STRMQMLSR command.

*Receiving on TCP:*

Receiving channel programs are started in response to a startup request from the sending channel. To respond to the startup request, a listener program has to be started to detect incoming network requests and start the associated channel. You start this listener program with the STRMQMLSR command.

You can start more than one listener for each queue manager. By default, the STRMQMLSR command uses port 1414 but you can override this value. To override the default setting, add the following statements to the qm.ini file of the selected queue manager. In this example, the listener is required to use port 2500:

```
TCP:
Port=2500
```

The qm.ini file is located in this IFS directory: /QIBM/UserData/mqm/qmgrs/ *queue manager name*.

This new value is read only when the TCP listener is started. If you have a listener already running, this change is not be seen by that program. To use the new value, stop the listener and issue the STRMQMLSR command again. Now, whenever you use the STRMQMLSR command, the listener defaults to the new port.

Alternatively, you can specify a different port number on the STRMQMLSR command. For example:

```
STRMQMLSR MQMNAME(queue manager name) PORT(2500)
```

This change makes the listener default to the new port for the duration of the listener job.

#### **Using the TCP SO\_KEEPALIVE option**

If you want to use the SO\_KEEPALIVE option (for more information, see “Checking that the other end of the channel is still available” on page 840 ) you must add the following entry to your queue manager configuration file (qm.ini in the IFS directory, /QIBM/UserData/mqm/qmgrs/ *queue manager name* ):

```
TCP:
KeepAlive=yes
```

You must then issue the following command:

```
CFGTCP
```

Select option 3 (Change TCP Attributes). You can now specify a time interval in minutes. You can specify a value in the range 1 through 40320 minutes; the default is 120.

#### **Using the TCP listener backlog option**

When receiving on TCP, a maximum number of outstanding connection requests is set. This number can be considered a *backlog* of requests waiting on the TCP port for the listener to accept the request.

The default listener backlog value on IBM i is 255. If the backlog reaches this value, the TCP connection is rejected and the channel is not able to start.

For MCA channels, this results in the channel going into a RETRY state and retrying the connection at a later time.

For client connections, the client receives an MQRC\_Q\_MGR\_NOT\_AVAILABLE reason code from MQCONN and can retry the connection at a later time.

However, to avoid this error, you can add an entry in the qm.ini file:

```
ListenerBacklog = n
```

This overrides the default maximum number of outstanding requests (255) for the TCP listener.

**Note:** Some operating systems support a larger value than the default. If necessary, this value can be used to avoid reaching the connection limit.

### Defining an LU 6.2 connection on IBM i:

Define the LU 6.2 communications details by using a mode name, TP name, and connection name of a fully qualified LU 6.2 connection.

The initiated end of the link must have a routing entry definition to complement this CSI object. Further information about managing work requests from remote LU 6.2 systems is available in the *IBM i Programming: Work Management Guide*.

See the *Multiplatform APPC Configuration Guide* and the following table for information.

*Table 133. Settings on the local IBM i system for a remote queue manager platform*

| Remote platform             | TPNAME                                                                                                                            |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| z/OS or MVS                 | The same as in the corresponding side information about the remote queue manager.                                                 |
| IBM i                       | The same as the compare value in the routing entry on the IBM i system.                                                           |
| HP Integrity NonStop Server | The same as the TPNAME specified in the receiver-channel definition.                                                              |
| UNIX and Linux systems      | The invocable Transaction Program defined in the remote LU 6.2 configuration.                                                     |
| Windows                     | As specified in the Windows Run Listener command, or the invocable Transaction Program that was defined using TpSetup on Windows. |

If you have more than one queue manager on the same computer, ensure that the TPnames in the channel definitions are unique.



## Related concepts:

“Initiating end (Sender)”

Use the CRTMQMCHL command to define a channel of transport type \*LU62.

“Initiated end (Receiver)” on page 900

Use the CRTMQMCHL command to define the receiving end of the message channel link with transport type \*LU62.

*Initiating end (Sender):*

Use the CRTMQMCHL command to define a channel of transport type \*LU62.

Use of the CSI object is optional in IBM MQ for IBM i V5.3 or later.

The initiating end panel is shown in Figure LU 6.2 communication setup panel - initiating end. To obtain the complete panel as shown, press F10 from the first panel.

```
Create Comm Side Information (CRTCSI)

Type choices, press Enter.

Side information > WINSDOA1 Name
Library > QSYS Name, *CURLIB
Remote location > WINSDOA1 Name
Transaction program > MQSERIES

Text 'description' *BLANK

Additional Parameters

Device *LOC Name, *LOC
Local location *LOC Name, *LOC, *NETATR
Mode JSTMOD92 Name, *NETATR
Remote network identifier . . . *LOC Name, *LOC, *NETATR, *NONE
Authority *LIBCRTAUT Name, *LIBCRTAUT, *CHANGE...

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
```

Figure 114. LU 6.2 communication setup panel - initiating end

Complete the initiating end fields as follows:

### Side information

Give this definition a name that is used to store the side information object to be created, for example, WINSDOA1.

**Note:** For LU 6.2, the link between the message channel definition and the communication connection is the **Connection name** field of the message channel definition at the sending end. This field contains the name of the CSI object.

### Library

The name of the library where this definition is stored.

The CSI object must be available in a library accessible to the program serving the message channel, for example, QSYS, QMQM, and QGPL.

If the name is incorrect, missing, or cannot be found then an error occurs on channel startup.

### Remote location

Specifies the remote location name with which your program communicates.

In short, this required parameter contains the logical unit name of the partner at the remote system, as defined in the device description that is used for the communication link between the two systems.

The **Remote location** name can be found by issuing the command DSPNETA on the remote system and seeing the default local location name.

### **Transaction program**

Specifies the name (up to 64 characters) of the transaction program on the remote system to be started. It can be a transaction process name, a program name, the channel name, or a character string that matches the **Compare value** in the routing entry.

This parameter is required.

**Note:** To specify SNA service transaction program names, enter the hexadecimal representation of the service transaction program name. For example, to specify a service transaction program name with a hexadecimal representation of 21F0F0F1, you would enter X'21F0F0F1'.

More information about SNA service transaction program names is in the *SNA Transaction Programmer's Reference* manual for LU Type 6.2.

If the receiving end is another IBM i system, the **Transaction program** name is used to match the CSI object at the sending end with the routing entry at the receiving end. This name must be unique for each queue manager on the target IBM i system. See the **Program to call** parameter under Initiated end (Receiver). See also the **Comparison data: compare value** parameter in the Add Routing Entry panel.

### **Text description**

A description (up to 50 characters) to remind you of the intended use of this connection.

### **Device**

Specifies the name of the device description used for the remote system. The possible values are:

**\*LOC** The device is determined by the system.

#### **Device-name**

Specify the name of the device that is associated with the remote location.

### **Local location**

Specifies the local location name. The possible values are:

**\*LOC** The local location name is determined by the system.

#### **\*NETATR**

The LCLLOCNAME value specified in the system network attributes is used.

#### **Local-location-name**

Specify the name of your location. Specify the local location if you want to indicate a specific location name for the remote location. The location name can be found by using the DSPNETA command.

**Mode** Specifies the mode used to control the session. This name is the same as the Common Programming Interface (CPI)- Communications Mode\_Name. The possible values are:

#### **\*NETATR**

The mode in the network attributes is used.

#### **BLANK**

Eight blank characters are used.

#### **Mode-name**

Specify a mode name for the remote location.

**Note:** Because the mode determines the transmission priority of the communications session, it might be useful to define different modes depending on the priority of the messages being sent; for example MQMODE\_HI, MQMODE\_MED, and MQMODE\_LOW. (You can have more than one CSI pointing to the same location.)

### **Remote network identifier**

Specifies the remote network identifier used with the remote location. The possible values are:

**\*LOC** The remote network ID for the remote location is used.

**\*NETATR**

The remote network identifier specified in the network attributes is used.

**\*NONE**

The remote network has no name.

### **Remote-network-id**

Specify a remote network ID. Use the DSPNETA command at the remote location to find the name of this network ID. It is the 'local network ID' at the remote location.

### **Authority**

Specifies the authority you are giving to users who do not have specific authority to the object, who are not on an authorization list, and with a group profile that has no specific authority to the object. The possible values are:

**\*LIBCRTAUT**

Public authority for the object is taken from the CRTAUT parameter of the specified library. This value is determined at create time. If the CRTAUT value for the library changes after the object is created, the new value does not affect existing objects.

**\*CHANGE**

Change authority allows the user to perform basic functions on the object, however, the user cannot change the object. Change authority provides object operational authority and all data authority.

**\*ALL** The user can perform all operations except those operations limited to the owner or controlled by authorization list management authority. The user can control the existence of the object and specify the security for the object, change the object, and perform basic functions on the object. The user can change ownership of the object.

**\*USE** Use authority provides object operational authority and read authority.

**\*EXCLUDE**

Exclude authority prevents the user from accessing the object.

### **Authorization-list**

Specify the name of the authorization list with authority that is used for the side information.

*Initiated end (Receiver):*

Use the CRTMQMCHL command to define the receiving end of the message channel link with transport type \*LU62.

Leave the CONNECTION NAME field blank and ensure that the corresponding details match the sending end of the channel. For details, see *Creating a channel*.

To enable the initiating end to start the receiving channel, add a routing entry to a subsystem at the initiated end. The subsystem must be the one that allocates the APPC device used in the LU 6.2 sessions. Therefore, it must have a valid communications entry for that device. The routing entry calls the program that starts the receiving end of the message channel.

Use the IBM i commands (for example, ADDRTGE) to define the end of the link that is initiated by a communication session.

The initiated end panel is shown in LU 6.2 communication setup panel - add routing entry.

```
Add Routing Entry (ADDRTGE)

Type choices, press Enter.

Subsystem description QCMN Name
Library *LIBL Name, *LIBL, *CURLIB
Routing entry sequence number . 1 1-9999
Comparison data:
Compare value MQSERIES

Starting position 37 1-80
Program to call AMQCRC6B Name, *RTGDTA
Library QMAS400 Name, *LIBL, *CURLIB
Class *SBSD Name, *SBSD
Library *LIBL Name, *LIBL, *CURLIB
Maximum active routing steps . . *NOMAX 0-1000, *NOMAX
Storage pool identifier 1 1-10

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
```

Figure 115. LU 6.2 communication setup panel - initiated end

### Subsystem description

The name of your subsystem where this definition resides. Use the IBM i WRKSBSD command to view and update the appropriate subsystem description for the routing entry.

### Routing entry sequence number

A unique number in your subsystem to identify this communication definition. You can use values in the range 1 - 9999.

### Comparison data: Compare value

A text string to compare with the string received when the session is started by a **Transaction program** parameter, as shown in Figure 1. The character string is derived from the Transaction program field of the sender CSI.

### Comparison data: Starting position

The character position in the string where the comparison is to start.

**Note:** The starting position field is the character position in the string for comparison, and this position is always 37.

## Program to call

The name of the program that runs the inbound message program to be called to start the session.

The program, AMQCRC6A, is called for the default queue manager. This program is supplied with IBM MQ for IBM i and sets up the environment and then calls AMQCRS6A.

For additional queue managers:

- Each queue manager has a specific LU 6.2 invocable program located in its library. This program is called AMQCRC6B and is automatically generated when the queue manager is created.
- Each queue manager requires a specific routing entry with unique routing data to be added. This routing data must match the **Transaction program** name supplied by the requesting system (see Initiating end (Sender) ).

An example is shown in LU 6.2 communication setup panel - display routing entries:

```
Display Routing Entries
System: MY400
Subsystem description: QCMN Status: ACTIVE

Type options, press Enter.
5=Display details

Start
Opt Seq Nbr Program Library Compare Value Pos
10 *RTGDTA 'QZSCSRVR' 37
20 *RTGDTA 'QZRCSRVR' 37
30 *RTGDTA 'QZHQTRG' 37
50 *RTGDTA 'QVPPRINT' 37
60 *RTGDTA 'QNPSRVR' 37
70 *RTGDTA 'QNMAPPINGD' 37
80 QNMAREXECD QSYS 'AREXECD' 37
90 AMQCRC6A QMQMBW 'MQSERIES' 37
100 *RTGDTA 'QTFDWNLD' 37
150 *RTGDTA 'QMFRFCVR' 37

F3=Exit F9=Display all detailed descriptions F12=Cancel
```

Figure 116. LU 6.2 communication setup panel - initiated end

In LU 6.2 communication setup panel - display routing entries, sequence number 90 represents the default queue manager and provides compatibility with configurations from previous releases (that is, V3R2, V3R6, V3R7, and V4R2) of IBM MQ for IBM i. These releases allow one queue manager only. Sequence numbers 92 and 94 represent two additional queue managers called ALPHA and BETA that are created with libraries QMALPHA and QMBETA.

**Note:** You can have more than one routing entry for each queue manager by using different routing data. These entries give the option of different job priorities depending on the classes used.

**Class** The name and library of the class used for the steps started through this routing entry. The class defines the attributes of the routing step's running environment and specifies the job priority. An appropriate class entry must be specified. Use, for example, the WRKCLS command to display existing classes or to create a class. Further information about managing work requests from remote LU 6.2 systems is available in the *IBM i Programming: Work Management Guide*.

## Note on Work Management

The AMQCRS6A job is not able to take advantage of the normal IBM i work management features that are documented in Work management because it is not started in the same way as other IBM MQ jobs. To change the runtime properties of the LU62 receiver jobs, you can make one of the following changes:

- Alter the class description that is specified on the routing entry for the AMQCRS6A job
- Change the job description on the communications entry

See the *IBM i Programming: Work Management Guide* for more information about configuring Communication Jobs.

## Configuring a queue manager cluster

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

### Before you begin

For an introduction to clustering concepts, see Clusters.

When you are designing your queue manager cluster you have to make some decisions. See Example clusters and Designing clusters.

#### Related tasks:

“Moving a cluster topic definition to a different queue manager” on page 1011

For either topic host routed or direct routed clusters, you might need to move a cluster topic definition when decommissioning a queue manager, or because a cluster queue manager has failed or is unavailable for a significant period of time.

#### Related information:

DELETE TOPIC

## Defining components of a cluster

Clusters are composed of queue managers, cluster channels, and cluster queues. You can define cluster queues, and modify some aspects of default cluster objects. You can get configuration and status information about auto-defined channels, and about the relationship between individual cluster-sender channels and transmission queues.

See the following subtopics for information about defining each of the cluster components:

#### Related concepts:

Components of a cluster

Cluster channels

#### Related tasks:

“Setting up a new cluster” on page 913

Follow these instructions to set up the example cluster. Separate instructions describe setting up the cluster on TCP/IP, LU 6.2, and with a single transmission queue or multiple transmission queues. Test the cluster works by sending a message from one queue manager to the other.

“Adding a queue manager to a cluster” on page 924

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using the single cluster transmission queue `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

#### Related information:

Defining cluster topics

## Defining cluster queues:

A cluster queue is a queue that is hosted by a cluster queue manager and made available to other queue managers in the cluster. Define a cluster queue as a local queue on the cluster queue manager where the queue is hosted. Specify the name of the cluster the queue belongs to.

The following example shows a **runmqsc** command to define a cluster queue with the **CLUSTER** option:

```
DEFINE QLOCAL(Q1) CLUSTER(SALES)
```

A cluster queue definition is advertised to other queue managers in the cluster. The other queue managers in the cluster can put messages to a cluster queue without needing a corresponding remote-queue definition. A cluster queue can be advertised in more than one cluster by using a cluster namelist.

When a queue is advertised, any queue manager in the cluster can put messages to it. To put a message, the queue manager must find out, from the full repositories, where the queue is hosted. Then it adds some routing information to the message and puts the message on a cluster transmission queue.

A cluster queue can be a queue that is shared by members of a queue-sharing group in IBM MQ for z/OS.

## Binding

You can create a cluster in which more than one queue manager hosts an instance of the same cluster queue. Make sure that all the messages in a sequence are sent to the same instance of the queue. You can bind a series of messages to a particular queue by using the **MQ00\_BIND\_ON\_OPEN** option on the **MQOPEN** call.

## Cluster transmission queues

A queue manager can store messages for other queue managers in a cluster on multiple transmission queues. You can configure a queue manager to store messages on multiple cluster transmission queues in two different ways. If you set the queue manager attribute **DEFCLXQ** to **CHANNEL**, a different cluster transmission queue is created automatically from **SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE** for each cluster-sender channel. If you set the **CLCHNAME** transmission queue option to match one or more cluster-senders channel, the queue manager can store messages for the matching channels on that transmission queue.

A message for a cluster queue on a different queue manager is placed upon a cluster transmission queue before being sent. A cluster-sender channel transfers the messages from a cluster transmission queue to cluster-receiver channels on other queue managers. By default, one system defined cluster transmission queue holds all the messages that are to be transferred to other cluster queue managers. The queue is called **SYSTEM.CLUSTER.TRANSMIT.QUEUE**. A queue manager that is part of a cluster can send messages on this cluster transmission queue to any other queue manager in the same cluster.

A definition for the single **SYSTEM.CLUSTER.TRANSMIT.QUEUE** queue is created by default on every queue manager except on z/OS. On z/OS, the definition can be defined with the supplied sample **CSQ4INSX**.

You can configure a queue manager to transfer messages to other clustered queue managers using multiple transmission queues. You can define additional cluster transmission queues manually, or have the queue manager create the queues automatically.

To have the queues created automatically by the queue manager, change the queue manager attribute **DEFCLXQ** from **SCTQ** to **CHANNEL**. The result is the queue manager creates an individual cluster transmission queue for each cluster-sender channel that is created. The transmission queues are created as permanent dynamic queues from the model queue, **SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE**. The name of each

permanent dynamic queue is `SYSTEM.CLUSTER.TRANSMIT.ChannelName`. The name of the cluster-sender channel that each permanent dynamic cluster transmit queue is associated with is set in the local transmission queue attribute `CLCHNAME`. Messages for remote clustered queue managers are placed on the permanent dynamic cluster transmission queue for the associated cluster-sender channel, rather than on `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

To create the cluster transmission queues manually, create a local queue with the `USAGE` attribute set to `XMITQ`, and the `CLCHNAME` attribute set to a generic channel name that resolves to one or more cluster-sender channels; see `ClusterChannelName`. If you create cluster transmission queues manually, you have the choice of associating the transmission queue with a single cluster-sender channel, or with multiple cluster-sender channels. The `CLCHNAME` attribute is a generic-name, which means you can place multiple wildcard characters, `"*"`, in the name.

Except for the initial cluster-sender channels that you create manually to connect a queue manager to a full repository, cluster-sender channels are created automatically. They are created automatically when there is a message to transfer to a cluster queue manager. They are created with the same name as the name of the cluster-receiver channel that receives cluster messages for that particular cluster on the destination queue manager.

If you follow a naming convention for cluster-receiver channels, it is possible to define a generic value for `CLCHNAME` that filters different kinds of cluster messages to different transmission queues. For example, if you follow the naming convention for cluster-receiver channels of `ClusterName.QmgrName`, then the generic name `ClusterName.*` filters messages for different clusters onto different transmission queues. You must define the transmission queues manually, and set `CLCHNAME` in each transmission queue to `ClusterName.*`.

Changes to the association of cluster transmission queues to cluster-sender channels do not take immediate effect. The currently associated transmission queue that a cluster-sender channel is servicing might contain messages that are in the process of being transferred by the cluster-sender channel. Only when no messages on the currently associated transmission queue are being processed by a cluster-sender channel, can the queue manager change the association of the cluster-sender channel a different transmission queue. This can occur either when no messages remain on the transmission queue to be processed by the cluster-sender channel, or when processing of messages is suspended and the cluster-sender channel has no "in-flight" messages. When this happens, any unprocessed messages for the cluster-sender channel are transferred to the newly associated transmission queue, and the association of the cluster-sender channel changes.

You can create a remote queue definition that resolves to a cluster transmission queue. In the definition, queue manager `QMX` is in the same cluster as the local queue manager, and there is no transmission queue, `QMX`.

```
DEFINE QREMOTE(A) RNAME(B) RQMNAME(QMX)
```

During queue name resolution, the cluster transmission queue takes precedence over the default transmission queue. A message put to `A` is stored on the cluster transmission queue and then sent to the remote queue `B` on `QMX`.

Queue managers can also communicate with other queue managers that are not part of a cluster. You must define channels and a transmission queue to the other queue manager, in the same way as in a distributed-queuing environment.

**Note:** Applications must write to queues that resolve to the cluster transmission queue, and must not write directly to the cluster transmission queue.



## Auto definition of remote queues

A queue manager in a cluster does not need a remote-queue definition for remote queues in the cluster. The cluster queue manager finds the location of a remote queue from the full repository. It adds routing information to the message and puts it on the cluster transmission queue. IBM MQ automatically creates a definition equivalent to a remote-queue definition so that the message can be sent.

You cannot alter or delete an automatically created remote-queue definition. However, by using the `DISPLAY QUEUE runmqsc` command with the `CLUSINFO` attribute, you can view all of the local queues on a queue manager as well as all of the cluster queues, including cluster queues on remote queue managers. For example:

```
DISPLAY QUEUE(*) CLUSINFO
```

### Related information:

Cluster queues

ClusterChannelName (MQCHAR20)

## Working with auto-defined cluster-sender channels:

After you introduce a queue manager to a cluster by making its initial `CLUSSDR` and `CLUSRCVR` definitions, IBM MQ automatically makes other cluster-sender channel definitions when they are needed. You can view information about auto-defined cluster-sender channels, but you cannot modify them. To modify their behavior, you can use a channel auto-definition exit.

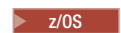
## Before you begin

For an introduction to auto-defined channels, see [Auto-defined cluster-sender channels](#).

## About this task

Auto-defined cluster-sender channels are created by the cluster as and when needed, and they remain active until they are shut down using the normal disconnect-interval rules.

On platforms other than z/OS, the OAM (object authority manager) is not aware of the existence of auto-defined cluster-sender channels. If you issue **start**, **stop**, **ping**, **reset**, or **resolve** commands on an auto-defined cluster-sender channel, the OAM checks to see whether you are authorized to perform the same action on the matching cluster-receiver channel.

 On z/OS, you can secure an auto-defined cluster-sender channel in the same way as any other channel.

## Procedure

- Display information about the auto-defined channels for a given cluster queue manager.

You cannot see automatically defined channels using the `DISPLAY CHANNEL runmqsc` command. To see the auto-defined channels use the following command:

```
DISPLAY CLUSQMGR(qMgrName)
```

- Display the status of the auto-defined channel for a given `CLUSRCVR`.

To display the status of the auto-defined `CLUSSDR` channel corresponding to a `CLUSRCVR` channel definition you created, use the following command:

```
DISPLAY CHSTATUS(channelname)
```

- Use a channel auto-definition exit to modify the behavior of an auto-defined channel.

You can use the IBM MQ channel auto-definition exit if you want to write a user exit program to customize a cluster-sender channel or cluster-receiver channel. For example, you can use the channel auto-definition exit in a cluster environment to make any of the following modifications:

- Tailor communications definitions, that is, SNA LU6.2 names.
- Add or remove other exits, for example, security exits.
- Change the names of channel exits.

The name of the CLUSSDR channel exit is auto-generated from the CLUSRCVR channel definition, and therefore might not be appropriate for your needs - especially if the two ends of the channel are on different platforms.

The format of exit names is different on different platforms. For example:

- **z/OS** On the z/OS platform, the format of the SCYEXIT (*security exit name*) parameter is SCYEXIT('SECEXIT')
- **Windows** On Windows platforms, the format of the SCYEXIT (*security exit name*) parameter is SCYEXIT(' *drive:\path\library* (secexit)')

**Note:** **z/OS** If there is no channel auto-definition exit, the z/OS queue manager derives the CLUSSDR channel exit name from the CLUSRCVR channel definition on the other end of the channel. To derive the z/OS exit name from a non-z/OS name, the following algorithm is used:

- Exit names on platforms other than z/OS are of the general form *path/library (function)*.
- If *function* is present, up to eight chars of that are used.
- Otherwise, up to eight chars of *library* are used.

For example:

- /var/mqm/exits/myExit.so(MsgExit) converts to MSGEXIT
- /var/mqm/exits/myExit converts to MYEXIT
- /var/mqm/exits/myExit.so(ExitLongName) converts to EXITLONG

- For queue managers earlier than IBM MQ Version 7, set the **PROPCTL** attribute to a value of NONE.

Each auto-defined cluster-sender channel is based on the corresponding cluster-receiver channel. Before IBM MQ Version 7, the cluster-receiver channel does not have a **PROPCTL** attribute, so this attribute is therefore set to COMPAT in the auto-defined cluster-sender channel.

If the cluster needs to use **PROPCTL** to remove application headers such as RFH2 from messages going from an IBM MQ Version 7 or later queue manager to a queue manager on an earlier version of IBM MQ, you must write a channel auto-definition exit that sets **PROPCTL** to a value of NONE.

- Use the channel attribute LOCLADDR. to control aspects of addressing.
  - To enable an outbound (TCP) channel to use a particular IP address, port or port range, use the channel attribute LOCLADDR. This is useful if you have more than one network card and you want a channel to use a specific one for outbound communications.
  - To specify a virtual IP address on CLUSSDR channels, use the IP address from the LOCLADDR on a manually defined CLUSSDR. To specify the port range, use the port range from the CLUSRCVR.
  - If a cluster needs to use LOCLADDR to get the outbound communication channels to bind to a specific IP address, you can write a channel auto-definition exit to force the LOCLADDR value into any of their automatically defined CLUSSDR channels. You must also specify it in the manually defined CLUSSDR channel.
  - Put a port number or port range in the LOCLADDR of a CLUSRCVR channel, if you want all the queue managers in a cluster to use a specific port or range of ports, for all their outbound communications.

**Note:** Do not put an IP address in the LOCLADDR field of a CLUSRCVR channel, unless all queue managers are on the same server. The LOCLADDR IP address is propagated to the auto-defined CLUSSDR channels of all queue managers that connect using the CLUSRCVR channel.

**distributed** On distributed platforms, it is possible to set a default local address value that will be used for all sender channels that do not have a local address defined. The default value is defined by setting the MQ\_LCLADDR environment variable prior to starting the queue manager. The format of the value matches that of MQSC attribute LOCLADDR.

**Related information:**

Local Address (LOCLADDR)

**Working with default cluster objects:**

You can alter the default channel definitions in the same way as any other channel definition, by running MQSC or PCF commands. Do not alter the default queue definitions, except for SYSTEM.CLUSTER.HISTORY.QUEUE.


For a full list of these objects, see Default cluster objects. The following list only includes those objects that you can change.

**SYSTEM.CLUSTER.HISTORY.QUEUE**

Each queue manager in a cluster has a local queue called SYSTEM.CLUSTER.HISTORY.QUEUE. SYSTEM.CLUSTER.HISTORY.QUEUE is used to store the history of cluster state information for service purposes.

In the default object settings, SYSTEM.CLUSTER.HISTORY.QUEUE is set to PUT ( ENABLED). To suppress history collection change the setting to PUT ( DISABLED).

**SYSTEM.CLUSTER.TRANSMIT.QUEUE**

Each queue manager has a definition for a local queue called SYSTEM.CLUSTER.TRANSMIT.QUEUE. SYSTEM.CLUSTER.TRANSMIT.QUEUE is the default transmission queue for all messages to all queues and queue managers that are within clusters. You can change the default transmission queue for each cluster-sender channel to SYSTEM.CLUSTER.TRANSMIT. *ChannelName*, by changing the queue manager attribute DEFXMITQ , except on z/OS . You cannot delete SYSTEM.CLUSTER.TRANSMIT.QUEUE. It is also used to define authorization checks whether the default transmission queue that is used is SYSTEM.CLUSTER.TRANSMIT.QUEUE or SYSTEM.CLUSTER.TRANSMIT. *ChannelName*.

**Related information:**

Default cluster objects

**Working with cluster transmission queues and cluster-sender channels:**

Messages between clustered queue managers are stored on cluster transmission queues and forwarded by cluster-sender channels. At any point in time, a cluster-sender channel is associated with one transmission queue. If you change the configuration of the channel, it might switch to a different transmission queue next time it starts. The processing of this switch is automated, and transactional.

Run the following MQSC command to display the transmission queues that cluster-sender channels are associated with:

```
DISPLAY CHSTATUS(*) WHERE(CHLTYPE EQ CLUSSDR)
AMQ8417: Display Channel Status details.
CHANNEL(TO.QM2) CHLTYPE(CLUSSDR)
CONNAME(9.146.163.190(1416)) CURRENT
RQMNAME(QM2) STATUS(STOPPED)
SUBSTATE() XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
```

The transmission queue shown in the saved channel status of a stopped cluster-sender channel might change when the channel starts again. “Selection of default transmission queues by cluster-sender channels” on page 908 describes the process of selecting a default transmission queue; “Selection of manually defined transmission queues by cluster-sender channels” on page 909 describes the process of selecting a manually defined transmission queue.

When any cluster-sender channel starts it rechecks its association with transmission queues. If the configuration of transmission queues, or the queue manager defaults, changes, it might reassociate the channel with a different transmission queue. If the channel restarts with a different transmission queue as

a result of a configuration change, a process of transferring messages to the newly associated transmission queue takes place. “How the process to switch cluster-sender channel to a different transmission queue works” on page 910 describes the process of transferring a cluster-sender channel from one transmission queue to another.

The behavior of cluster-sender channels is different to sender and server channels. They remain associated with the same transmission queue until the channel attribute **XMITQ** is altered. If you alter the transmission queue attribute on a sender or server channel, and restart it, messages are not transferred from the old transmission queue to the new one.

Another difference between cluster-sender channels, and sender or server channels, is that multiple cluster-sender channels can open a cluster transmission queue, but only one sender or server channel can open a normal transmission queue. Until Version 7.5, cluster connections shared the single cluster transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. From Version 7.5 onwards, you have the option of cluster-sender channels not sharing transmission queues. Exclusivity is not enforced; it is an outcome of the configuration. You can configure the path a message takes in a cluster so that it does not share any transmission queues or channels with messages that flow between other applications. See *Clustering: Planning how to configure cluster transmission queues and “Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager”* on page 957.

To configure a cluster-sender channel to use a transmission queue other than `SYSTEM.CLUSTER.TRANSMIT.QUEUE` on z/OS, you need to enable version 8 new function, using the mode of operation ( `OPMODE` ) system parameter in the `CSQ6SYSP` macro.

### **Selection of default transmission queues by cluster-sender channels**

A cluster transmission queue is either a system default queue, with a name that starts with `SYSTEM.CLUSTER.TRANSMIT`, or a manually defined queue. A cluster-sender channel is associated with a cluster transmission queue in one of two ways: by the default cluster transmission queue mechanism, or by manual configuration.

The default cluster transmission queue is set as a queue manager attribute, **DEFCLXQ**. Its value is either `SCTQ` or `CHANNEL`. New and migrated queue managers are set to `SCTQ`. You can alter the value to `CHANNEL`.

If `SCTQ` is set, the default cluster transmission queue is `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. Every cluster-sender channel can open this queue. The cluster-sender channels that do open the queue are the ones that are not associated with manually defined cluster transmission queues.

If `CHANNEL` is set, then the queue manager can create a separate permanent dynamic transmission queue for every cluster-sender channel. Each queue is named `SYSTEM.CLUSTER.TRANSMIT.ChannelName` and is created from the model queue, `SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE`. Each cluster-sender channel that is not associated with a manually defined cluster transmission queue is associated with a permanent-dynamic cluster transmission queue. The queue is created by the queue manager when it requires a separate cluster transmission queue for the cluster destination served by this cluster-sender channel, and no queue exists.

Some cluster destinations can be served by cluster-sender channels associated with manually defined transmission queues, and others by the default queue or queues. In the association of cluster-sender channels with transmission queues, the manually defined transmission queues always take precedence over the default transmission queues.

The precedence of cluster transmission queues is illustrated in Figure 117 on page 909. The only cluster-sender channel not associated with a manually defined cluster transmission queue is `CS.QM1`. It is not associated with a manually defined transmission queue, because none of the channel names in the **CLCHNAME** attribute of the transmission queues match `CS.QM1`.

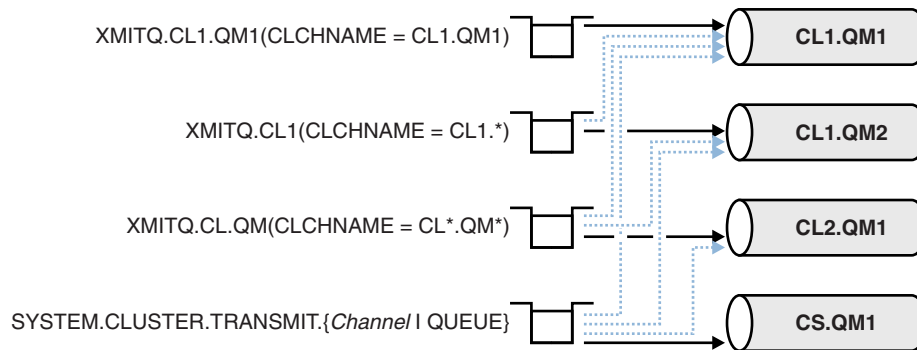


Figure 117. Transmission queue / cluster-sender channel precedence

### Selection of manually defined transmission queues by cluster-sender channels

A manually defined queue has the transmission queue attribute **USAGE** attribute set to XMITQ, and the cluster channel name attribute **CLCHNAME** set to a specific or generic channel name.

If the name in the **CLCHNAME** queue attribute matches a cluster-sender channel name, the channel is associated with the queue. The name is either an exact match, if the name contains no wildcards, or it the best match, if the name contains wildcards.

If **CLCHNAME** definitions on multiple transmission queues match the same cluster-sender channel, the definitions are said to overlap. To resolve the ambiguity there is an order of precedence between matches. Exact matches always take precedence. Figure 117 shows associations between transmission queues and cluster-sender channels. The black arrows show actual associations, and the gray arrows, potential associations. The precedence order of transmission queues in Figure 117 is,

#### XMITQ.CL1.QM1

The transmission queue XMITQ.CL1.QM1 has its **CLCHNAME** attribute set to CL1.QM1. The definition of the **CLCHNAME** attribute, CL1.QM1, has no wildcards, and takes precedence over any other **CLCHNAME** attributes, defined on other transmission queues, that match with wildcards. The queue manager stores any cluster message that is to be transferred by the CL1.QM1 cluster-sender channel on the XMITQ.CL1.QM1 transmission queue. The only exception is if multiple transmission queues have their **CLCHNAME** attribute set to CL1.QM1. In that case, the queue manager stores messages for the CL1.QM1 cluster-sender channel on any of those queues. It selects a queue arbitrarily when the channel starts. It might select a different queue when the channel starts again.

#### XMITQ.CL1

The transmission queue XMITQ.CL1 has its **CLCHNAME** attribute set to CL1.\*. The definition of the **CLCHNAME** attribute, CL1.\*, has one trailing wildcard, which matches the name of any cluster-sender channel that starts with CL1.. The queue manager stores any cluster message that is to be transferred by any cluster-sender channel whose name begins with CL1. on the transmission queue XMITQ.CL1, unless there is a transmission queue with a more specific match, such as the queue XMITQ.CL1.QM1. One trailing wildcard makes the definition less specific than a definition with no wildcards, and more specific than a definition with multiple wildcards, or wildcards that are followed by more trailing characters.

#### XMITQ.CL.QM

XMITQ.CL.QM is the name of the transmission queue with its **CLCHNAME** attribute set to CL\*.QM\*. The definition of CL\*.QM\* has two wildcards, which match the name of any cluster-sender channel that starts with CL., and either includes or ends with QM. The match is less specific than a match with one wildcard.

## **SYSTEM.CLUSTER.TRANSMIT.** *channelName* |**QUEUE**

If no transmission queue has a **CLCHNAME** attribute that matches the name of the cluster-sender channel that the queue manager is to use, then the queue manager uses the default cluster transmission queue. The default cluster transmission queue is either the single system cluster transmission queue, **SYSTEM.CLUSTER.TRANSMIT.QUEUE**, or a system cluster transmission queue that the queue manager created for a specific cluster-sender channel, **SYSTEM.CLUSTER.TRANSMIT.** *channelName*. Which queue is the default depends on the setting of the queue manager **DEFXMITQ** attribute.

**Tip:** Unless you have a clear need for overlapping definitions, avoid them as they can lead to complicated configurations that are hard to understand.

### **How the process to switch cluster-sender channel to a different transmission queue works**

To change the association of cluster-sender channels with cluster transmission queues, change the **CLCHNAME** parameter of any transmission queue or the queue manager parameter **DEFCLXQ** at any time. Nothing happens immediately. Changes occur only when a channel starts. When it starts, it checks whether to continue forwarding messages from the same transmission queue. Three kinds of change alter the association of a cluster-sender channel with a transmission queue.

1. Redefining the **CLCHNAME** parameter of the transmission queue the cluster-sender channel is currently associated with to be less specific or blank, or deleting the cluster transmission queue when the channel is stopped.

Some other cluster transmission queue might now be a better match for the channel name. Or, if no other transmission queues match the name of the cluster-sender channel, the association must revert to the default transmission queue.

2. Redefining the **CLCHNAME** parameter of any other cluster transmission queue, or adding a cluster transmission queue.

The **CLCHNAME** parameter of another transmission queue might now be a better match for the cluster-sender channel than the transmission queue the cluster-sender channel is currently associated with. If the cluster-sender channel is currently associated with a default cluster transmission queue, it might become associated with a manually defined cluster transmission queue.

3. If the cluster-sender channel is currently associated with a default cluster transmission queue, changing the **DEFCLXQ** queue manager parameter.

If the association of a cluster-sender channel changes, when the channel starts it switches its association to the new transmission queue. During the switch, it ensures that no messages are lost. Messages are transferred to the new transmission queue in the order in which the channel would transfer the messages to the remote queue manager.

**Remember:** In common with any forwarding of messages in a cluster, you must put messages into groups to ensure that messages that must be delivered in order are delivered in order. On rare occasions, messages can get out of order in a cluster.

The switch process goes through the following transactional steps. If the switch process is interrupted, the current transactional step is resumed when the channel restarts again.

#### **Step 1 - Process messages from the original transmission queue**

The cluster-sender channel is associated with the new transmission queue, which it might share with other cluster-sender channels. Messages for the cluster-sender channel continue to be placed on the original transmission queue. A transitional switch process transfers messages from the original transmission queue onto the new transmission queue. The cluster-sender channel forwards the messages from the new transmission queue to the cluster-receiver channel. The channel status shows the cluster-sender channel still associated with the old transmission queue.

The switch process continues to transfer newly arrived messages as well. This step continues until the number of remaining messages to be forwarded by the switch process reaches zero. When the number of messages reaches zero, the procedure moves onto step 2.

During step 1, disk activity for the channel increases. Persistent messages are committed off the first transmission queue and onto the second transmission queue. This disk activity is in addition to messages being committed when they are placed on and removed from the transmission queue as part of transferring the messages normally. Ideally, no messages arrive during the switching process, so the transition can take place as quickly as possible. If messages do arrive, they are processed by the switch process.

### **Step 2 - Process messages from the new transmission queue**

As soon as no messages remain on the original transmission queue for the cluster-sender channel, new messages are placed directly on the new transmission queue. The channel status shows the cluster-sender channel is associated with the new transmission queue. The following message is written to the queue manager error log: " AMQ7341 The transmission queue for channel *ChannelName* is *QueueName* ."

### **Multiple cluster transmission queues and cluster transmission queue attributes**

You have a choice of forwarding cluster messages to different queue managers storing the messages on a single cluster transmission queue, or multiple queues. With one queue, you have one set of cluster transmission queue attributes to set and query; with multiple queues, you have multiple sets. For some attributes, having multiple sets is an advantage: for example querying queue depth tells you how many messages are waiting to be forwarded by one or a set of channels, rather than by all channels. For other attributes, having multiple sets is a disadvantage: for example, you probably do not want to configure the same access permissions for every cluster transmission queue. For this reason, access permissions are always checked against the profile for `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, and not against profiles for a particular cluster transmission queue. If you want to apply more granular security checks, see Access control and multiple cluster transmission queues.

### **Multiple cluster-sender channels and multiple transmission queues**

A queue manager stores a message on a cluster transmission queue before forwarding it on a cluster-sender channel. It selects a cluster-sender channel that is connected to the destination for the message. It might have a choice of cluster-sender channels that all connect to the same destination. The destination might be the same physical queue, connected by multiple cluster-sender channels to a single queue manager. The destination might also be many physical queues with the same queue name, hosted on different queue managers in the same cluster. Where there is a choice of cluster-sender channels connected to a destination, the workload balancing algorithm chooses one. The choice depends on a number of factors; see The cluster workload management algorithm.

In Figure 118 on page 912, `CL1.QM1`, `CL1.QM2` and `CS.QM1` are all channels that might lead to the same destination. For example, if you define `Q1` in `CL1` on `QM1` and `QM2` then `CL1.QM1` and `CL1.QM2` both provide routes to the same destination, `Q1`, on two different queue managers. If the channel `CS.QM1` is also in `CL1`, it too is a channel that a message for `Q1` can take. The cluster membership of `CS.QM1` might be defined by a cluster namelist, which is why the channel name does not include a cluster name in its construction. Depending on the workload balancing parameters, and the sending application, some messages for `Q1` might be placed on each of the transmission queues, `XMITQ.CL1.QM1`, `XMITQ.CL1` and `SYSTEM.CLUSTER.TRANSMIT.CS.QM1`.

If you intend to separate out message traffic, so that messages for the same destination do not share queues or channels with messages for different destinations, you must consider how to divide traffic onto different cluster-sender channels first, and then how to separate messages for a particular channel onto a different transmission queue. Cluster queues on the same cluster, on the same queue manager, normally share the same cluster channels. Defining multiple cluster transmission queues alone is not sufficient to

separate cluster message traffic onto different queues. Unless you separate messages for different destination queues onto different channels, the messages share the same cluster transmission queue.

A straightforward way to separate the channels that messages take, is to create multiple clusters. On any queue manager in each cluster, define only one cluster queue. Then, if you define a different cluster-receiver channel for each cluster/queue manager combination, the messages for each cluster queue do not share a cluster channel with messages for other cluster queues. If you define separate transmission queues for the cluster channels, the sending queue manager stores messages for only one cluster queue on each transmission queue. For example, if you want two cluster queues not to share resources, you can either place them in different clusters on the same queue manager, or on different queue managers in the same cluster.

The choice of cluster transmission queue does not affect the workload balancing algorithm. The workload balancing algorithm chooses which cluster-sender channel to forward a message. It places the message on the transmission queue that is serviced by that channel. If the workload balancing algorithm is called on to choose again, for instance if the channel stops, it might be able to select a different channel to forward the message. If it does choose a different channel, and the new channel forwards messages from a different cluster transmission queue, the workload balancing algorithm transfers the message to the other transmission queue.

In Figure 118, two cluster-sender channels, CS.QM1 and CS.QM2, are associated with the default system transmission queue. When the workload balancing algorithm stores a message on SYSTEM.CLUSTER.TRANSMIT.QUEUE, or any other cluster transmission queue, the name of the cluster-sender channel that is to forward the message is stored in the correlation ID of the message. Each channel forwards just those messages that match the correlation ID with the channel name.

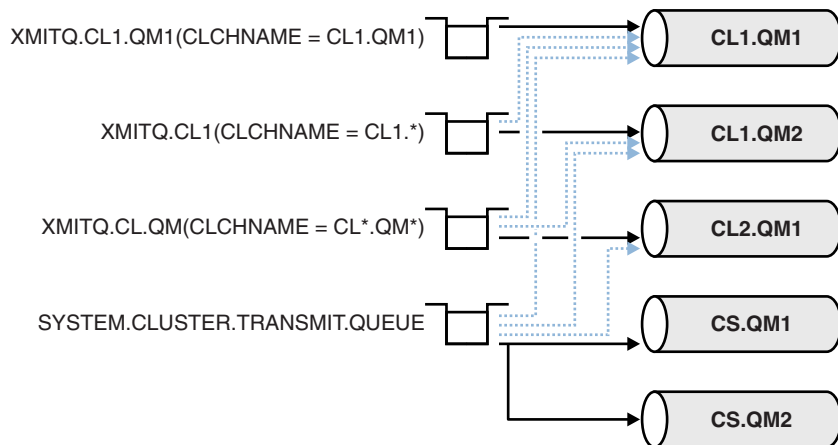


Figure 118. Multiple cluster sender channels

If CS.QM1 stops, the messages on the transmission queue for that cluster-sender channel are examined. Those messages that can be forwarded by another channel are reprocessed by the workload balancing algorithm. Their correlation ID is reset to an alternative cluster-sender channel name. If the alternative cluster-sender channel is CS.QM2, the message remains on SYSTEM.CLUSTER.TRANSMIT.QUEUE. If the alternative channel is CL1.QM1, the workload balancing algorithm transfers the message to XMITQ.CL1.QM1. When the cluster-sender channel restarts, new messages, and messages that were not flagged for a different cluster-sender channel, are transferred by the channel again.

You might change the association between transmission queues and cluster-sender channels on a running system. You might change a CLCHNAME parameter on a transmission queue, or, change the DEFCLXQ queue manager parameter. When a channel that is affected by the change restarts, it starts the transmission queue switching process; see “How the process to switch cluster-sender channel to a different transmission queue works” on page 910.



The process to switch the transmission queue starts when the channel is restarted. The workload rebalancing process starts when the channel is stopped. The two process can run in parallel.

The simple case is when stopping a cluster-sender channel does not cause the rebalancing process to change the cluster-sender channel that is to forward any messages on the queue. This case is when no other cluster-sender channel can forward the messages to the correct destination. With no alternative cluster-sender channel to forward the messages to their destination, the messages remain flagged for the same cluster-sender channel after the cluster-sender channel stops. When the channel starts, if a switch is pending, the switching processes moves the messages to a different transmission queue where they are processed by the same cluster-sender channel.

The more complex case is where more than one cluster-sender channel can process some messages to the same destination. You stop and restart the cluster-sender channel to trigger the transmission queue switch. In many cases, by the time you restart the channel, the workload balancing algorithm has already moved messages from the original transmission queue to different transmission queues served by different cluster-sender channels. Only those messages that cannot be forwarded by a different cluster-sender channel remain to be transferred to the new transmission queue. In some cases, if the channel is restarted quickly, some messages that could be transferred by the workload balancing algorithm remain. In which case some remaining messages are switched by the workload balancing process, and some by the process of switching the transmission queue.

**Related concepts:**

“Calculating the size of the log” on page 1136  
Estimating the size of log a queue manager needs.

**Related tasks:**

“Creating two-overlapping clusters with a gateway queue manager” on page 947  
Follow the instructions in the task to construct overlapping clusters with a gateway queue manager. Use the clusters as a starting point for the following examples of isolating messages to one application from messages to other applications in a cluster.

“Adding a queue manager to a cluster: separate transmission queues” on page 926  
Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

“Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954

Modify the configuration of overlapping clusters that use a gateway queue manager. After the modification messages are transferred to an application from the gateway queue manager without using the same transmission queue or channels as other cluster messages. The solution uses an additional cluster transmission queue to separate message traffic to a single queue manager in a cluster.

“Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 957

Modify the configuration of overlapping clusters that use a gateway queue manager. After the modification messages are transferred to an application from the gateway queue manager without using the same transmission queue or channels as other cluster messages. The solution uses an additional cluster to isolate the messages to a particular cluster queue.

**Related information:**

Cluster channels

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

## Setting up a new cluster

Follow these instructions to set up the example cluster. Separate instructions describe setting up the cluster on TCP/IP, LU 6.2, and with a single transmission queue or multiple transmission queues. Test the cluster works by sending a message from one queue manager to the other.

## Before you begin

- Instead of following these instructions, you can use one of the wizards supplied with MQ Explorer to create a cluster like the one created by this task. Right-click the Queue Manager Clusters folder, then click **New > Queue Manager Cluster**, and follow the instructions given in the wizard.
- For background information to aid your understanding of the steps taken to set up a cluster, see “Defining cluster queues” on page 903, Cluster channels and Listeners.

## About this task

You are setting up a new IBM MQ network for a chain store. The store has two branches, one in London and one in New York. The data and applications for each store are hosted by systems running separate queue managers. The two queue managers are called LONDON and NEWYORK. The inventory application runs on the system in New York, connected to queue manager NEWYORK. The application is driven by the arrival of messages on the INVENTQ queue, hosted by NEWYORK. The two queue managers, LONDON and NEWYORK, are to be linked in a cluster called INVENTORY so that they can both put messages to the INVENTQ.

Figure 119 shows what this cluster looks like.

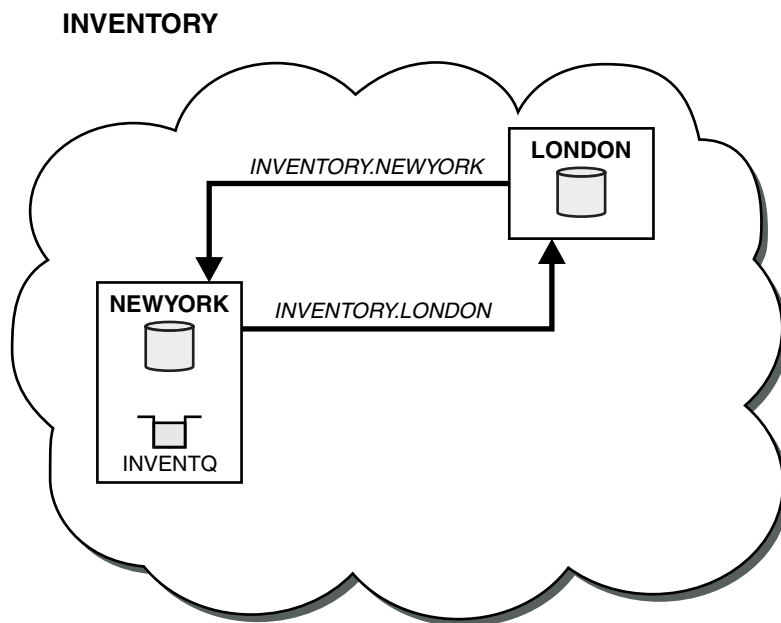


Figure 119. The INVENTORY cluster with two queue managers

You can configure each queue manager in the cluster to send messages to other queue managers in the cluster using different cluster transmission queues.

The instructions to set up the cluster vary a little by transport protocol, number of transmission queues, or platform. You have a choice of three combinations. The verification procedure remains the same for all combinations.

## Procedure

- “Setting up a cluster using TCP/IP with a single transmission queue per queue manager” on page 915
- “Setting up a cluster on TCP/IP using multiple transmission queues per queue manager” on page 917
- “Setting up a cluster using LU 6.2 on z/OS” on page 920
- “Verifying the cluster” on page 922

## Results

Figure 119 on page 914 shows the INVENTORY cluster setup by this task.

Clearly, INVENTORY is a small cluster. However, it is useful as a proof of concept. The important thing to understand about this cluster is the scope it offers for future enhancement.

### Related tasks:

“Configuring a queue manager cluster” on page 902

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

### Related information:

Clusters

Comparison of clustering and distributed queuing


Components of a cluster

### Setting up a cluster using TCP/IP with a single transmission queue per queue manager:

#### Before you begin

- The queue manager attribute, **DEFCLXQ**, must be left as its default value, SCTQ.

#### About this task

Follow these steps to set up a cluster on AIX, HP-UX, IBM i, Linux, Solaris, and Windows using the transport protocol TCP/IP.  On z/OS, you must follow the instructions in “Defining a TCP connection on z/OS” on page 1284 to set up the TCP/IP connection, rather than defining the listeners in step 4 on page 916. Otherwise, the steps are the same for z/OS, but error messages are written to the console, rather than to the queue manager error log.

#### Procedure

1. Decide on the organization of the cluster and its name.

You decided to link the two queue managers, LONDON and NEWYORK, into a cluster. A cluster with only two queue managers offers only marginal benefit over a network that is to use distributed queuing. It is a good way to start and it provides scope for future expansion. When you open new branches of your store, you are able to add the new queue managers to the cluster easily. Adding new queue managers does not disrupt the existing network; see “Adding a queue manager to a cluster” on page 924.

For the time being, the only application you are running is the inventory application. The cluster name is INVENTORY.

2. Decide which queue managers are to hold full repositories.

In any cluster you must nominate at least one queue manager, or preferably two, to hold full repositories. In this example, there are only two queue managers, LONDON and NEWYORK, both of which hold full repositories.

- a. You can perform the remaining steps in any order.
- b. As you proceed through the steps, warning messages might be written to the queue-manager log. The messages are a result of missing definitions that you have yet to add.

Examples of the responses to the commands are shown in a box like this after each step in this task. These examples show the responses returned by IBM MQ for AIX. The responses vary on other platforms.

- c. Before proceeding with these steps, make sure that the queue managers are started.
3. Alter the queue-manager definitions to add repository definitions.  
On each queue manager that is to hold a full repository, alter the local queue-manager definition, using the ALTER QMGR command and specifying the REPOS attribute:  
ALTER QMGR REPOS(INVENTORY)

```
1 : ALTER QMGR REPOS(INVENTORY)
AMQ8005: Websphere MQ queue manager changed.
```

For example, if you enter:

- a. runmqsc LONDON
- b. ALTER QMGR REPOS(INVENTORY)

LONDON is changed to a full repository.

4. Define the listeners.  
Define a listener that accepts network requests from other queue managers for every queue manager in the cluster. On the LONDON queue managers, issue the following command:  
DEFINE LISTENER(LONDON\_LS) TRPTYPE(TCP) CONTROL(QMGR)

The CONTROL attribute ensures that the listener starts and stops when the queue manager does. The listener is not started when it is defined, so it must be manually started the first time with the following MQSC command:  
START LISTENER(LONDON\_LS)

Issue similar commands for all the other queue managers in the cluster, changing the listener name for each one.

There are several ways to define these listeners, as shown in Listeners.

5. Define the CLUSRCVR channel for the LONDON queue manager.  
On every queue manager in a cluster, you define a cluster-receiver channel on which the queue manager can receive messages. See Cluster-receiver channel: CLUSRCVR . The CLUSRCVR channel defines the connection name of the queue manager. The connection name is stored in the repositories, where other queue managers can refer to it. The CLUSTER keyword shows the availability of the queue manager to receive messages from other queue managers in the cluster.  
In this example the channel name is INVENTORY.LONDON, and the connection name (CONNAME) is the network address of the machine the queue manager resides on, which is LONDON.CHSTORE.COM. The network address can be entered as an alphanumeric DNS host name, or an IP address in either in IPv4 dotted decimal form. For example, 192.0.2.0, or IPv6 hexadecimal form; for example 2001:DB8:0204:acff:fe97:2c34:fde0:3485. The port number is not specified, so the default port (1414) is used.  
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)  
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)  
DESCR('TCP Cluster-receiver channel for queue manager LONDON')

```

1 : DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-receiver channel for queue manager LONDON')
AMQ8014: Websphere MQ channel created.
07/09/98 12:56:35 No repositories for cluster 'INVENTORY'

```

6. Define the CLUSRCVR channel for the NEWYORK queue manager. If the channel listener is using the default port, typically 1414, and the cluster does not include a queue manager on z/OS, you can omit the CONNAME

```

DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSRCVR) TRPTYPE(TCP) CLUSTER(INVENTORY)
DESCR('TCP Cluster-receiver channel for queue manager NEWYORK')

```

7. Define the CLUSSDR channel on the LONDON queue manager.

You manually define a CLUSSDR channel from every full repository queue manager to every other full repository queue manager in the cluster. See Cluster-sender channel: CLUSSDR . In this case, there are only two queue managers, both of which hold full repositories. They each need a manually-defined CLUSSDR channel that points to the CLUSRCVR channel defined at the other queue manager. The channel names given on the CLUSSDR definitions must match the channel names on the corresponding CLUSRCVR definitions. When a queue manager has definitions for both a cluster-receiver channel and a cluster-sender channel in the same cluster, the cluster-sender channel is started.

```

DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from LONDON to repository at NEWYORK')

```

```

1 : DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from LONDON to repository at NEWYORK')
AMQ8014: Websphere MQ channel created.
07/09/98 13:00:18 Channel program started.

```

8. Define the CLUSSDR channel on the NEWYORK queue manager.

```

DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from NEWYORK to repository at LONDON')

```

9. Define the cluster queue INVENTQ

Define the INVENTQ queue on the NEWYORK queue manager, specifying the CLUSTER keyword.

```

DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)

```

```

1 : DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
AMQ8006: Websphere MQ queue created.

```

The CLUSTER keyword causes the queue to be advertised to the cluster. As soon as the queue is defined it becomes available to the other queue managers in the cluster. They can send messages to it without having to make a remote-queue definition for it.

All the definitions are complete. On all platforms, start a listener program on each queue manager. The listener program waits for incoming network requests and starts the cluster-receiver channel when it is needed.

**Setting up a cluster on TCP/IP using multiple transmission queues per queue manager:**

## About this task

Follow these steps to set up a cluster on AIX, HP-UX, IBM i, Linux, Solaris, and Windows using the transport protocol TCP/IP. The repository queue managers are configured to use a different cluster transmission queue to send messages to each other, and to other queue managers in the cluster. If you add queue managers to the cluster that are also to use different transmission queues, follow the task, “Adding a queue manager to a cluster: separate transmission queues” on page 926.

## Procedure

1. Decide on the organization of the cluster and its name.

You decided to link the two queue managers, LONDON and NEWYORK, into a cluster. A cluster with only two queue managers offers only marginal benefit over a network that is to use distributed queuing. It is a good way to start and it provides scope for future expansion. When you open new branches of your store, you are able to add the new queue managers to the cluster easily. Adding new queue managers does not disrupt the existing network; see “Adding a queue manager to a cluster” on page 924.

For the time being, the only application you are running is the inventory application. The cluster name is INVENTORY.

2. Decide which queue managers are to hold full repositories.

In any cluster you must nominate at least one queue manager, or preferably two, to hold full repositories. In this example, there are only two queue managers, LONDON and NEWYORK, both of which hold full repositories.

- a. You can perform the remaining steps in any order.
- b. As you proceed through the steps, warning messages might be written to the queue-manager log. The messages are a result of missing definitions that you have yet to add.

Examples of the responses to the commands are shown in a box like this after each step in this task. These examples show the responses returned by IBM MQ for AIX. The responses vary on other platforms.

- c. Before proceeding with these steps, make sure that the queue managers are started.

3. Alter the queue-manager definitions to add repository definitions.

On each queue manager that is to hold a full repository, alter the local queue-manager definition, using the ALTER QMGR command and specifying the REPOS attribute:

```
ALTER QMGR REPOS(INVENTORY)
```

```
1 : ALTER QMGR REPOS(INVENTORY)
AMQ8005: Websphere MQ queue manager changed.
```

For example, if you enter:

- a. runmqsc LONDON
- b. ALTER QMGR REPOS(INVENTORY)

LONDON is changed to a full repository.

4. Alter the queue-manager definitions to create separate cluster transmission queues for each destination.

```
ALTER QMGR DEFCLXQ(CHANNEL)
```

On each queue manager that you add to the cluster decide whether to use separate transmission queues or not. See the topics, “Adding a queue manager to a cluster” on page 924 and “Adding a queue manager to a cluster: separate transmission queues” on page 926.

5. Define the listeners.

Define a listener that accepts network requests from other queue managers for every queue manager in the cluster. On the LONDON queue managers, issue the following command:

```
DEFINE LISTENER(LONDON_LS) TRPTYPE(TCP) CONTROL(QMGR)
```

The CONTROL attribute ensures that the listener starts and stops when the queue manager does.

The listener is not started when it is defined, so it must be manually started the first time with the following MQSC command:

```
START LISTENER(LONDON_LS)
```

Issue similar commands for all the other queue managers in the cluster, changing the listener name for each one.

There are several ways to define these listeners, as shown in Listeners.

6. Define the CLUSRCVR channel for the LONDON queue manager.

On every queue manager in a cluster, you define a cluster-receiver channel on which the queue manager can receive messages. See Cluster-receiver channel: CLUSRCVR . The CLUSRCVR channel defines the connection name of the queue manager. The connection name is stored in the repositories, where other queue managers can refer to it. The CLUSTER keyword shows the availability of the queue manager to receive messages from other queue managers in the cluster.

In this example the channel name is INVENTORY.LONDON, and the connection name (CONNAME) is the network address of the machine the queue manager resides on, which is LONDON.CHSTORE.COM. The network address can be entered as an alphanumeric DNS host name, or an IP address in either in IPv4 dotted decimal form. For example, 192.0.2.0, or IPv6 hexadecimal form; for example 2001:DB8:0204:acff:fe97:2c34:fde0:3485. The port number is not specified, so the default port (1414) is used.

```
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-receiver channel for queue manager LONDON')
```

```
1 : DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-receiver channel for queue manager LONDON')
AMQ8014: Websphere MQ channel created.
07/09/98 12:56:35 No repositories for cluster 'INVENTORY'
```

7. Define the CLUSRCVR channel for the NEWYORK queue manager. If the channel listener is using the default port, typically 1414, and the cluster does not include a queue manager on z/OS, you can omit the CONNAME

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSRCVR) TRPTYPE(TCP) CLUSTER(INVENTORY)
DESCR('TCP Cluster-receiver channel for queue manager NEWYORK')
```

8. Define the CLUSSDR channel on the LONDON queue manager.

You manually define a CLUSSDR channel from every full repository queue manager to every other full repository queue manager in the cluster. See Cluster-sender channel: CLUSSDR . In this case, there are only two queue managers, both of which hold full repositories. They each need a manually-defined CLUSSDR channel that points to the CLUSRCVR channel defined at the other queue manager. The channel names given on the CLUSSDR definitions must match the channel names on the corresponding CLUSRCVR definitions. When a queue manager has definitions for both a cluster-receiver channel and a cluster-sender channel in the same cluster, the cluster-sender channel is started.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from LONDON to repository at NEWYORK')
```

```
1 : DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from LONDON to repository at NEWYORK')
AMQ8014: Websphere MQ channel created.
07/09/98 13:00:18 Channel program started.
```

9. Define the CLUSSDR channel on the NEWYORK queue manager.

```
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('TCP Cluster-sender channel from NEWYORK to repository at LONDON')
```

10. Define the cluster queue INVENTQ

Define the INVENTQ queue on the NEWYORK queue manager, specifying the CLUSTER keyword.

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
```

```
1 : DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
AMQ8006: Websphere MQ queue created.
```

The CLUSTER keyword causes the queue to be advertised to the cluster. As soon as the queue is defined it becomes available to the other queue managers in the cluster. They can send messages to it without having to make a remote-queue definition for it.

All the definitions are complete. On all platforms, start a listener program on each queue manager. The listener program waits for incoming network requests and starts the cluster-receiver channel when it is needed.

## Setting up a cluster using LU 6.2 on z/OS:

### Procedure

1. Decide on the organization of the cluster and its name.

You decided to link the two queue managers, LONDON and NEWYORK, into a cluster. A cluster with only two queue managers offers only marginal benefit over a network that is to use distributed queuing. It is a good way to start and it provides scope for future expansion. When you open new branches of your store, you are able to add the new queue managers to the cluster easily. Adding new queue managers does not disrupt the existing network; see "Adding a queue manager to a cluster" on page 924.

For the time being, the only application you are running is the inventory application. The cluster name is INVENTORY.

2. Decide which queue managers are to hold full repositories.

In any cluster you must nominate at least one queue manager, or preferably two, to hold full repositories. In this example, there are only two queue managers, LONDON and NEWYORK, both of which hold full repositories.

- a. You can perform the remaining steps in any order.
- b. As you proceed through the steps, warning messages might be written the z/OS system console. The messages are a result of missing definitions that you have yet to add.
- c. Before proceeding with these steps, make sure that the queue managers are started.

3. Alter the queue-manager definitions to add repository definitions.



On each queue manager that is to hold a full repository, alter the local queue-manager definition, using the ALTER QMGR command and specifying the REPOS attribute:

```
ALTER QMGR REPOS(INVENTORY)
```

```
1 : ALTER QMGR REPOS(INVENTORY)
AMQ8005: Websphere MQ queue manager changed.
```

For example, if you enter:

- a. runmqsc LONDON
- b. ALTER QMGR REPOS(INVENTORY)

LONDON is changed to a full repository.

#### 4. Define the listeners.

 See The channel initiator on z/OS and “Receiving on LU 6.2” on page 1288.

The listener is not started when it is defined, so it must be manually started the first time with the following MQSC command:

```
START LISTENER(LONDON_LS)
```

Issue similar commands for all the other queue managers in the cluster, changing the listener name for each one.

#### 5. Define the CLUSRCVR channel for the LONDON queue manager.

On every queue manager in a cluster, you define a cluster-receiver channel on which the queue manager can receive messages. See Cluster-receiver channel: CLUSRCVR . The CLUSRCVR channel defines the connection name of the queue manager. The connection name is stored in the repositories, where other queue managers can refer to it. The CLUSTER keyword shows the availability of the queue manager to receive messages from other queue managers in the cluster.

```
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(LU62)
CONNNAME(LONDON.LUNAME) CLUSTER(INVENTORY)
MODENAME('#INTER') TPNAME('MQSERIES')
DESCR('LU62 Cluster-receiver channel for queue manager LONDON')
```

```
1 : DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) TRPTYPE(LU62)
CONNNAME(LONDON.LUNAME) CLUSTER(INVENTORY)
MODENAME('#INTER') TPNAME('MQSERIES')
DESCR('LU62 Cluster-receiver channel for queue manager LONDON')
AMQ8014: Websphere MQ channel created.
07/09/98 12:56:35 No repositories for cluster 'INVENTORY'
```

#### 6. Define the CLUSRCVR channel for the NEWYORK queue manager.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSRCVR) TRPTYPE(LU62)
CONNNAME(NEWYORK.LUNAME) CLUSTER(INVENTORY)
MODENAME('#INTER') TPNAME('MQSERIES')
DESCR('LU62 Cluster-receiver channel for queue manager NEWYORK')
```

#### 7. Define the CLUSSDR channel on the LONDON queue manager.

You manually define a CLUSSDR channel from every full repository queue manager to every other full repository queue manager in the cluster. See Cluster-sender channel: CLUSSDR . In this case, there are only two queue managers, both of which hold full repositories. They each need a manually-defined CLUSSDR channel that points to the CLUSRCVR channel defined at the other queue manager. The channel names given on the CLUSSDR definitions must match the channel names on the corresponding CLUSRCVR definitions. When a queue manager has definitions for both a cluster-receiver channel and a cluster-sender channel in the same cluster, the cluster-sender channel is started.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(LU62)
CONNAME(CPIC) CLUSTER(INVENTORY)
DESCR('LU62 Cluster-sender channel from LONDON to repository at NEWYORK')
```

```
1 : DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(LU62)
CONNAME(NEWYORK.LUNAME) CLUSTER(INVENTORY)
MODENAME('#INTER') TPNAME('MQSERIES')
DESCR('LU62 Cluster-sender channel from LONDON to repository at NEWYORK')
AMQ8014: Websphere MQ channel created.
07/09/98 13:00:18 Channel program started.
```

8. Define the CLUSSDR channel on the NEWYORK queue manager.

```
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSSDR) TRPTYPE(LU62)
CONNAME(LONDON.LUNAME) CLUSTER(INVENTORY)
DESCR('LU62 Cluster-sender channel from NEWYORK to repository at LONDON')
```

9. Define the cluster queue INVENTQ

Define the INVENTQ queue on the NEWYORK queue manager, specifying the CLUSTER keyword.

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
```

```
1 : DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
AMQ8006: Websphere MQ queue created.
```

The CLUSTER keyword causes the queue to be advertised to the cluster. As soon as the queue is defined it becomes available to the other queue managers in the cluster. They can send messages to it without having to make a remote-queue definition for it.

All the definitions are complete. On all platforms, start a listener program on each queue manager. The listener program waits for incoming network requests and starts the cluster-receiver channel when it is needed.

*Verifying the cluster:*

### About this task

You can verify the cluster in one or more of these ways:

1. Running administrative commands to display cluster and channel attributes.
2. Run the sample programs to send and receive messages on a cluster queue.
3. Write your own programs to send a request message to a cluster queue and reply with a response messages to an non-clustered reply queue.

### Procedure

1. Issue DISPLAY **runmqsc** commands to verify the cluster. The responses you see ought to be like the responses in the steps that follow.
  - a. From the NEWYORK queue manager, run the **DISPLAY CLUSQMGR** command:

```
dis clusqmgr(*)
```

```

1 : dis clusqmgr(*)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(NEWYORK) CLUSTER(INVENTORY)
CHANNEL(INVENTORY.NEWYORK)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(LONDON) CLUSTER(INVENTORY)
CHANNEL(INVENTORY.LONDON)

```

- b. From the NEWYORK queue manager, run the **DISPLAY CHANNEL STATUS** command:

```
dis chstatus(*)
```

```

1 : dis chstatus(*)
AMQ8417: Display Channel Status details.
CHANNEL(INVENTORY.NEWYORK) XMITQ()
CONNAME(192.0.2.0) CURRENT
CHLTYPE(CLUSRCVR) STATUS(RUNNING)
RQMNAME(LONDON)
AMQ8417: Display Channel Status details.
CHANNEL(INVENTORY.LONDON) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE) XMITQ(SYSTEM.CLUSTER.TRANSMIT.INVENTORY.LONDON)
CONNAME(192.0.2.1) CURRENT
CHLTYPE(CLUSSDR) STATUS(RUNNING)
RQMNAME(LONDON)

```

2. Send messages between the two queue managers, using **amqsput**.
  - a. On LONDON run the command **amqsput INVENTQ LONDON**.  
Type some messages, followed by a blank line.
  - b. On NEWYORK run the command **amqsget INVENTQ NEWYORK**.  
You now see the messages you entered on LONDON. After 15 seconds the program ends.
3. Send messages between the two queue managers using your own programs. In the following steps, LONDON puts a message to the INVENTQ at NEWYORK and receives a reply on its queue LONDON\_reply.
  - a. On LONDON put a messages to the cluster queue.
    - 1) Define a local queue called LONDON\_reply
    - 2) Set the MQOPEN options to MQ00\_OUTPUT
    - 3) Issue the MQOPEN call to open the queue INVENTQ
    - 4) Set the *ReplyToQ* name in the message descriptor to LONDON\_reply
    - 5) Issue the MQPUT call to put the message
    - 6) Commit the message
  - b. On NEWYORK receive the message on the cluster queue and put a reply to the reply queue.
    - 1) Set the MQOPEN options to MQ00\_BROWSE
    - 2) Issue the MQOPEN call to open the queue INVENTQ
    - 3) Issue the MQGET call to get the message from INVENTQ
    - 4) Retrieve the *ReplyToQ* name from the message descriptor
    - 5) Put the *ReplyToQ* name in the ObjectName field of the object descriptor
    - 6) Set the MQOPEN options to MQ00\_OUTPUT
    - 7) Issue the MQOPEN call to open LONDON\_reply at queue manager LONDON
    - 8) Issue the MQPUT call to put the message to LONDON\_reply
  - c. On LONDON receive the reply.
    - 1) Set the MQOPEN options to MQ00\_BROWSE
    - 2) Issue the MQOPEN call to open the queue LONDON\_reply

- 3) Issue the MQGET call to get the message from LONDON\_reply

## Adding a queue manager to a cluster

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using the single cluster transmission queue SYSTEM.CLUSTER.TRANSMIT.QUEUE.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster is set up as described in “Setting up a new cluster” on page 913. It contains two queue managers, LONDON and NEWYORK, which both hold full repositories.
- The queue manager PARIS is owned by the primary installation. If it is not, you must run the **setmqenv** command to set up the command environment for the installation that PARIS belongs to.
- TCP connectivity exists between all three systems, and the queue manager is configured with a TCP listener that starts under the control of the queue manager.

### About this task

1. A new branch of the chain store is being set up in Paris and you want to add a queue manager called PARIS to the cluster.
2. Queue manager PARIS sends inventory updates to the application running on the system in New York by putting messages on the INVENTQ queue.

Follow these steps to add a queue manager to a cluster.

### Procedure

1. Decide which full repository PARIS refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories. It gathers information about the cluster from a full repository and so builds up its own partial repository. Choose either of the repositories as the full repository. As soon as a new queue manager is added to the cluster it immediately learns about the other repository as well. Information about changes to a queue manager is sent directly to two repositories. In this example, you link PARIS to the queue manager LONDON, purely for geographical reasons.


**Note:** Perform the remaining steps in any order, after queue manager PARIS is started.

2. Define a CLUSRCVR channel on queue manager PARIS.

Every queue manager in a cluster must define a cluster-receiver channel on which it can receive messages. On PARIS, define:

```
DEFINE CHANNEL(INVENTORY.PARIS) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(PARIS.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for queue manager PARIS')
```

The cluster-receiver channel advertises the availability of the queue manager to receive messages from other queue managers in the cluster INVENTORY. Do not make definitions on other queue managers for a sending end to the cluster-receiver channel INVENTORY.PARIS. Other definitions are made automatically when needed. See Cluster channels.

3.  Start the channel initiator on IBM MQ for z/OS.
4. Define a CLUSSDR channel on queue manager PARIS.

When you add to a cluster a queue manager that is not a full repository, you define just one cluster-sender channel to make an initial connection to a full repository. See Cluster-sender channel: CLUSSDR . On PARIS, make the following definition for a CLUSSDR channel called INVENTORY.LONDON to the queue manager with the network address LONDON.CHSTORE.COM.

```

DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from PARIS to repository at LONDON')

```

5. Optional: If you are adding to a cluster a queue manager that has previously been removed from the same cluster, check that it is now showing as a cluster member. If not, complete the following extra steps:

- a. Issue the **REFRESH CLUSTER** command on the queue manager you are adding. This step stops the cluster channels, and gives your local cluster cache a fresh set of sequence numbers that are assured to be up-to-date within the rest of the cluster.

```
REFRESH CLUSTER(INVENTORY) REPOS(YES)
```

**Note:** For large clusters, using the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.

- b. Restart the CLUSSDR channel (for example, using the START CHANNEL command).
- c. Restart the CLUSRCVR channel.

## Results

The following figure shows the cluster set up by this task.

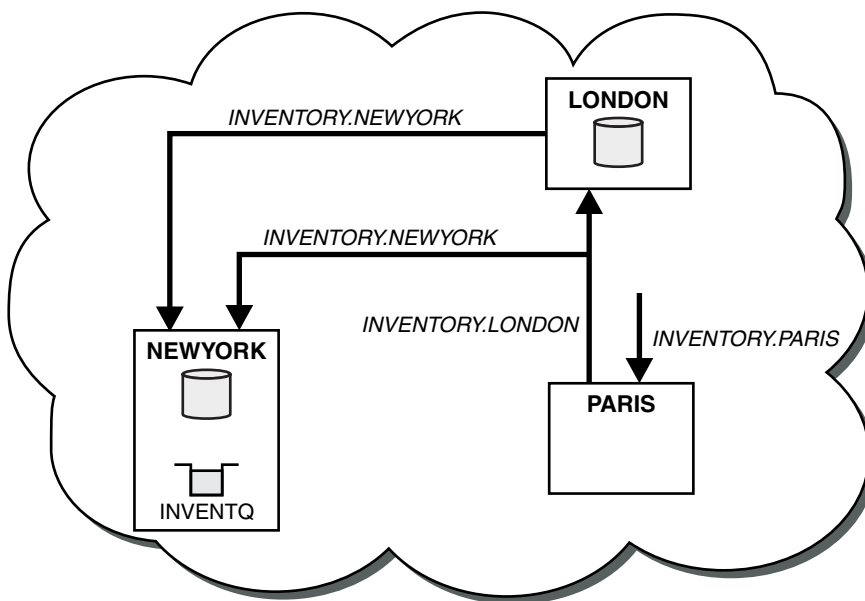


Figure 120. The INVENTORY cluster with three queue managers

By making only two definitions, a CLUSRCVR definition and a CLUSSDR definition, we added the queue manager PARIS to the cluster.

Now the PARIS queue manager learns, from the full repository at LONDON, that the INVENTQ queue is hosted by queue manager NEWYORK. When an application hosted by the system in Paris tries to put messages to the INVENTQ, PARIS automatically defines a cluster-sender channel to connect to the cluster-receiver channel INVENTORY.NEWYORK. The application can receive responses when its queue-manager name is specified as the target queue manager and a reply-to queue is provided.

## Adding a queue manager to a cluster: separate transmission queues:

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

### Before you begin

- The queue manager is not a member of any clusters.
- The cluster exists; there is a full repository to which this queue manager can connect directly and the repository is available. For the steps to create the cluster, see “Setting up a new cluster” on page 913.

### About this task

This task is an alternative to “Adding a queue manager to a cluster” on page 924, in which you add a queue manager to a cluster that places cluster messages on a single transmission queue.

In this task, you add a queue manager to a cluster that automatically creates separate cluster transmission queues for each cluster-sender channel.

To keep the number of definitions of queues small, the default is to use a single transmission queue. Using separate transmission queues is advantageous if you want to monitor traffic destined to different queue managers and different clusters. You might also want to separate traffic to different destinations to achieve isolation or performance goals.

### Procedure

1. Alter the default cluster channel transmission queue type.

Alter the queue manager PARIS:

```
ALTER QMGR DEFCLXQ(CHANNEL)
```

Every time the queue manager creates a cluster-sender channel to send a message to a queue manager, it creates a cluster transmission queue. The transmission queue is used only by this cluster-sender channel. The transmission queue is permanent-dynamic. It is created from the model queue, SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE, with the name SYSTEM.CLUSTER.TRANSMIT.  
*ChannelName*.

2. Decide which full repository PARIS refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories. It gathers information about the cluster from a full repository and so builds up its own partial repository. Choose either of the repositories as the full repository. As soon as a new queue manager is added to the cluster it immediately learns about the other repository as well. Information about changes to a queue manager is sent directly to two repositories. In this example, you link PARIS to the queue manager LONDON, purely for geographical reasons.

**Note:** Perform the remaining steps in any order, after queue manager PARIS is started.

3. Define a CLUSRCVR channel on queue manager PARIS.

Every queue manager in a cluster must define a cluster-receiver channel on which it can receive messages. On PARIS, define:

```
DEFINE CHANNEL(INVENTORY.PARIS) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(PARIS.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for queue manager PARIS')
```

The cluster-receiver channel advertises the availability of the queue manager to receive messages from other queue managers in the cluster INVENTORY. Do not make definitions on other queue managers for a sending end to the cluster-receiver channel INVENTORY.PARIS. Other definitions are made automatically when needed. See Cluster channels.

4. Define a CLUSSDR channel on queue manager PARIS.

When you add to a cluster a queue manager that is not a full repository, you define just one cluster-sender channel to make an initial connection to a full repository. See Cluster-sender channel: CLUSSDR . On PARIS, make the following definition for a CLUSSDR channel called INVENTORY.LONDON to the queue manager with the network address LONDON.CHSTORE.COM.

```
DEFINE CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from PARIS to repository at LONDON')
```

The queue manager automatically creates the permanent dynamic cluster transmission queue SYSTEM.CLUSTER.TRANSMIT.INVENTORY.LONDON from the model queue SYSTEM.CLUSTER.TRANSMIT.MODEL.QUEUE. It sets the CLCHNAME attribute of the transmission queue to INVENTORY.LONDON.

## Results

The following figure shows the cluster set up by this task.

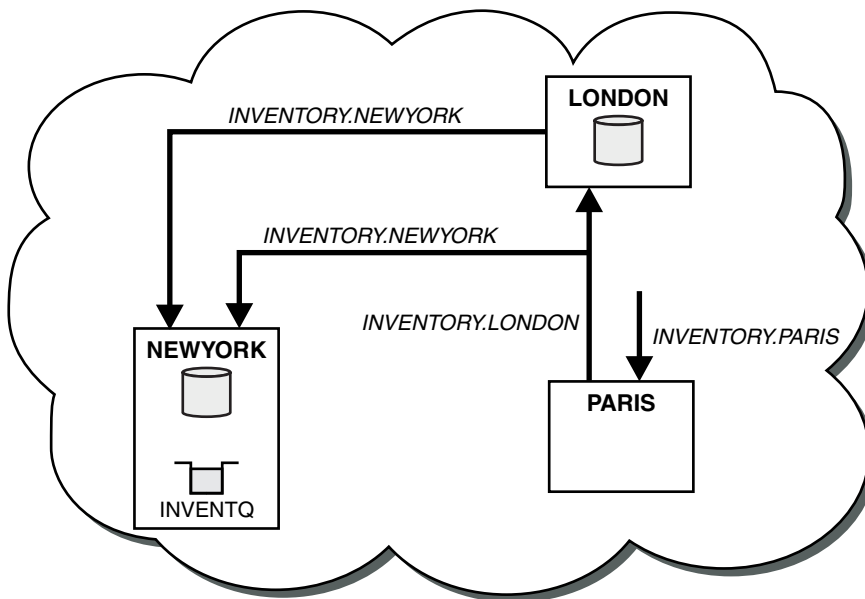


Figure 121. The INVENTORY cluster with three queue managers

By making only two definitions, a CLUSRCVR definition and a CLUSSDR definition, we added the queue manager PARIS to the cluster.

Now the PARIS queue manager learns, from the full repository at LONDON, that the INVENTQ queue is hosted by queue manager NEWYORK. When an application hosted by the system in Paris tries to put messages to the INVENTQ, PARIS automatically defines a cluster-sender channel to connect to the cluster-receiver channel INVENTORY.NEWYORK. The application can receive responses when its queue-manager name is specified as the target queue manager and a reply-to queue is provided.

## Adding a queue manager to a cluster by using DHCP:

Add a queue manager to a cluster, using DHCP. The task demonstrates omitting CONNAME value on a CLUSRCVR definition.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

The task demonstrates two special features:

- The ability to omit the CONNAME value on a CLUSRCVR definition.
- The ability to use +QMNAME+ on a CLUSSDR definition.

Neither feature is provided on z/OS.

Scenario:

- The INVENTORY cluster has been set up as described in “Setting up a new cluster” on page 913. It contains two queue managers, LONDON and NEWYORK, which both hold full repositories.
- A new branch of the chain store is being set up in Paris and you want to add a queue manager called PARIS to the cluster.
- Queue manager PARIS sends inventory updates to the application running on the system in New York by putting messages on the INVENTQ queue.
- Network connectivity exists between all three systems.
- The network protocol is TCP.
- The PARIS queue manager system uses DHCP, which means that the IP addresses might change on system restart.
- The channels between the PARIS and LONDON systems are named according to a defined naming convention. The convention uses the queue manager name of the full repository queue manager on LONDON.
- Administrators of the PARIS queue manager have no information about the name of the queue manager on the LONDON repository. The name of the queue manager on the LONDON repository is subject to change.

### About this task

Follow these steps to add a queue manager to a cluster by using DHCP.

#### Procedure

1. Decide which full repository PARIS refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories. It gathers information about the cluster from a full repository and so builds up its own partial repository. Choose either of the repositories as the full repository. As soon as a new queue manager is added to the cluster it immediately learns about the other repository as well. Information about changes to a queue manager is sent directly to two repositories. In this example we choose to link PARIS to the queue manager LONDON, purely for geographical reasons.

**Note:** Perform the remaining steps in any order, after queue manager PARIS is started.

2. Define a CLUSRCVR channel on queue manager PARIS.

Every queue manager in a cluster needs to define a cluster-receiver channel on which it can receive messages. On PARIS, define:



```

DEFINE CHANNEL(INVENTORY.PARIS) CHLTYPE(CLUSRCVR)
TRPTYPE(TCP) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for queue manager PARIS')

```

The cluster-receiver channel advertises the availability of the queue manager to receive messages from other queue managers in the cluster INVENTORY. You do not need to specify the CONNAME on the cluster-receiver channel. You can request IBM MQ to find out the connection name from the system, either by omitting CONNAME, or by specifying CONNAME(' '). IBM MQ generates the CONNAME value using the current IP address of the system; see CONNAME . There is no need to make definitions on other queue managers for a sending end to the cluster-receiver channel INVENTORY.PARIS. Other definitions are made automatically when needed.

3. Define a CLUSSDR channel on queue manager PARIS.

Every queue manager in a cluster needs to define one cluster-sender channel on which it can send messages to its initial full repository. On PARIS, make the following definition for a channel called INVENTORY.+QMNAME+ to the queue manager with the network address LONDON.CHSTORE.COM.

```

DEFINE CHANNEL(INVENTORY.+QMNAME+) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(LONDON.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from PARIS to repository at LONDON')

```

4. Optional: If you are adding to a cluster a queue manager that has previously been removed from the same cluster, check that it is now showing as a cluster member. If not, complete the following extra steps:

- a. Issue the **REFRESH CLUSTER** command on the queue manager you are adding. This step stops the cluster channels, and gives your local cluster cache a fresh set of sequence numbers that are assured to be up-to-date within the rest of the cluster.

```
REFRESH CLUSTER(INVENTORY) REPOS(YES)
```

**Note:** For large clusters, using the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.

- b. Restart the CLUSSDR channel (for example, using the START CHANNEL command).
- c. Restart the CLUSRCVR channel.

## Results

The cluster set up by this task is the same as for “Adding a queue manager to a cluster” on page 924:




Figure 122. The INVENTORY cluster with three queue managers

By making only two definitions, a CLUSRCVR definition and a CLUSSDR definition, we have added the queue manager PARIS to the cluster.

On the PARIS queue manager, the CLUSSDR containing the string +QMNAME+ starts. On the LONDON system IBM MQ resolves the +QMNAME+ to the queue manager name ( LONDON). IBM MQ then matches the definition for a channel called INVENTORY.LONDON to the corresponding CLUSRCVR definition.

IBM MQ sends back the resolved channel name to the PARIS queue manager. At PARIS, the CLUSSDR channel definition for the channel called INVENTORY.+QMNAME+ is replaced by an internally generated CLUSSDR definition for INVENTORY.LONDON. This definition contains the resolved channel name, but otherwise is the same as the +QMNAME+ definition that you made. The cluster repositories are also brought up to date with the channel definition with the newly resolved channel name.

**Note:**

1. The channel created with the +QMNAME+ name becomes inactive immediately. It is never used to transmit data.
2. Channel exits might see the channel name change between one invocation and the next.

Now the PARIS queue manager learns, from the repository at LONDON, that the INVENTQ queue is hosted by queue manager NEWYORK. When an application hosted by the system in Paris tries to put messages to the INVENTQ, PARIS automatically defines a cluster-sender channel to connect to the cluster-receiver channel INVENTORY.NEWYORK. The application can receive responses when its queue-manager name is specified as the target queue manager and a reply-to queue is provided.

## Related information:

DEFINE CHANNEL

## Adding a queue manager that hosts a queue

Add another queue manager to the cluster, to host another INVENTQ queue. Requests are sent alternately to the queues on each queue manager. No changes need to be made to the existing INVENTQ host.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager to a cluster” on page 924. It contains three queue managers; LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository. The inventory application runs on the system in New York, connected to the NEWYORK queue manager. The application is driven by the arrival of messages on the INVENTQ queue.
- A new store is being set up in Toronto. To provide additional capacity you want to run the inventory application on the system in Toronto as well as New York.
- Network connectivity exists between all four systems.
- The network protocol is TCP.

**Note:** The queue manager TORONTO contains only a partial repository. If you want to add a full-repository queue manager to a cluster, refer to “Moving a full repository to another queue manager” on page 935.

### About this task

Follow these steps to add a queue manager that hosts a queue.

### Procedure

1. Decide which full repository TORONTO refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories. It gathers information about the cluster from a full repository and so builds up its own partial repository. It is of no particular significance which repository you choose. In this example, we choose NEWYORK. Once the new queue manager has joined the cluster it communicates with both of the repositories.

2. Define the CLUSRCVR channel.

Every queue manager in a cluster needs to define a cluster-receiver channel on which it can receive messages. On TORONTO, define a CLUSRCVR channel:

```
DEFINE CHANNEL(INVENTORY.TORONTO) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(TORONTO.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for TORONTO')
```

The TORONTO queue manager advertises its availability to receive messages from other queue managers in the INVENTORY cluster using its cluster-receiver channel.

3. Define a CLUSSDR channel on queue manager TORONTO.

Every queue manager in a cluster needs to define one cluster-sender channel on which it can send messages to its first full repository. In this case choose NEWYORK. TORONTO needs the following definition:

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from TORONTO to repository at NEWYORK')
```

4. Optional: If you are adding to a cluster a queue manager that has previously been removed from the same cluster, check that it is now showing as a cluster member. If not, complete the following extra steps:

- a. Issue the **REFRESH CLUSTER** command on the queue manager you are adding. This step stops the cluster channels, and gives your local cluster cache a fresh set of sequence numbers that are assured to be up-to-date within the rest of the cluster.

```
REFRESH CLUSTER(INVENTORY) REPOS(YES)
```

**Note:** For large clusters, using the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.

- b. Restart the CLUSSDR channel (for example, using the START CHANNEL command).
  - c. Restart the CLUSRCVR channel.
5. Review the inventory application for message affinities.  
Before proceeding, ensure that the inventory application does not have any dependencies on the sequence of processing of messages and install the application on the system in Toronto.
  6. Define the cluster queue INVENTQ.

The INVENTQ queue, which is already hosted by the NEWYORK queue manager, is also to be hosted by TORONTO. Define it on the TORONTO queue manager as follows:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
```

## Results

Figure 123 shows the INVENTORY cluster set up by this task.




Figure 123. The INVENTORY cluster with four queue managers

The INVENTQ queue and the inventory application are now hosted on two queue managers in the cluster. This increases their availability, speeds up throughput of messages, and allows the workload to be distributed between the two queue managers. Messages put to INVENTQ by either TORONTO or NEWYORK are handled by the instance on the local queue manager whenever possible. Messages put by LONDON or PARIS are routed alternately to TORONTO or NEWYORK, so that the workload is balanced.

This modification to the cluster was accomplished without you having to alter the definitions on queue managers NEWYORK, LONDON, and PARIS. The full repositories in these queue managers are updated automatically with the information they need to be able to send messages to INVENTQ at TORONTO. The inventory application continues to function if one of the NEWYORK or the TORONTO queue manager becomes unavailable, and it has sufficient capacity. The inventory application must be able to work correctly if it is hosted in both locations.

As you can see from the result of this task, you can have the same application running on more than one queue manager. You can clustering to distribution workload evenly.

An application might not be able to process records in both locations. For example, suppose that you decide to add a customer-account query and update application running in LONDON and NEWYORK. An account record can only be held in one place. You could decide to control the distribution of requests by using a data partitioning technique. You can split the distribution of the records. You could arrange for half the records, for example for account numbers 00000 - 49999, to be held in LONDON. The other half, in the range 50000 - 99999, are held in NEWYORK. You could then write a cluster workload exit program to examine the account field in all messages, and route the messages to the appropriate queue manager.

## What to do next

Now that you have completed all the definitions, if you have not already done so start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program on queue manager TORONTO. The listener program waits for incoming network requests and starts the cluster-receiver channel when it is needed.

## Adding a queue-sharing group to existing clusters

Add a queue-sharing group on z/OS to existing clusters.

### Before you begin

#### Note:

1. For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.
2. Queue-sharing groups are supported only on IBM MQ for z/OS. This task is not applicable to other platforms.

#### Scenario:

- The INVENTORY cluster has been set up as described in “Setting up a new cluster” on page 913. It contains two queue managers, LONDON and NEWYORK.
- You want to add a queue-sharing group to this cluster. The group, QSGP, comprises three queue managers, P1, P2, and P3. They share an instance of the INVENTQ queue, which is to be defined by P1.

### About this task

Follow these steps to add new queue managers that host a shared queue.

### Procedure

1. Decide which full repository the queue managers refer to first.  
Every queue manager in a cluster must refer to one or other of the full repositories. It gathers information about the cluster from a full repository and so builds up its own partial repository. It is of no particular significance which full repository you choose. In this example, choose NEWYORK. Once the queue-sharing group has joined the cluster, it communicates with both of the full repositories.
2. Define the CLUSRCVR channels.

Every queue manager in a cluster needs to define a cluster-receiver channel on which it can receive messages. On P1, P2, and P3, define:

```
DEFINE CHANNEL(INVENTORY.Pn) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(Pn.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for sharing queue manager')
```

The cluster-receiver channel advertises the availability of each queue manager to receive messages from other queue managers in the cluster INVENTORY.

3. Define a CLUSSDR channel for the queue-sharing group.

Every member of a cluster needs to define one cluster-sender channel on which it can send messages to its first full repository. In this case we have chosen NEWYORK. One of the queue managers in the queue-sharing group needs the following group definition. The definition ensures that every queue manager has a cluster-sender channel definition.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY) QSGDISP(GROUP)
DESCR('Cluster-sender channel to repository at NEWYORK')
```

4. Define the shared queue.

Define the queue INVENTQ on P1 as follows:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY) QSGDISP(SHARED) CFSTRUCT(STRUCTURE)
```

Start the channel initiator and a listener program on the new queue manager. The listener program listens for incoming network requests and starts the cluster-receiver channel when it is needed.

## Results

Figure 124 shows the cluster set up by this task.

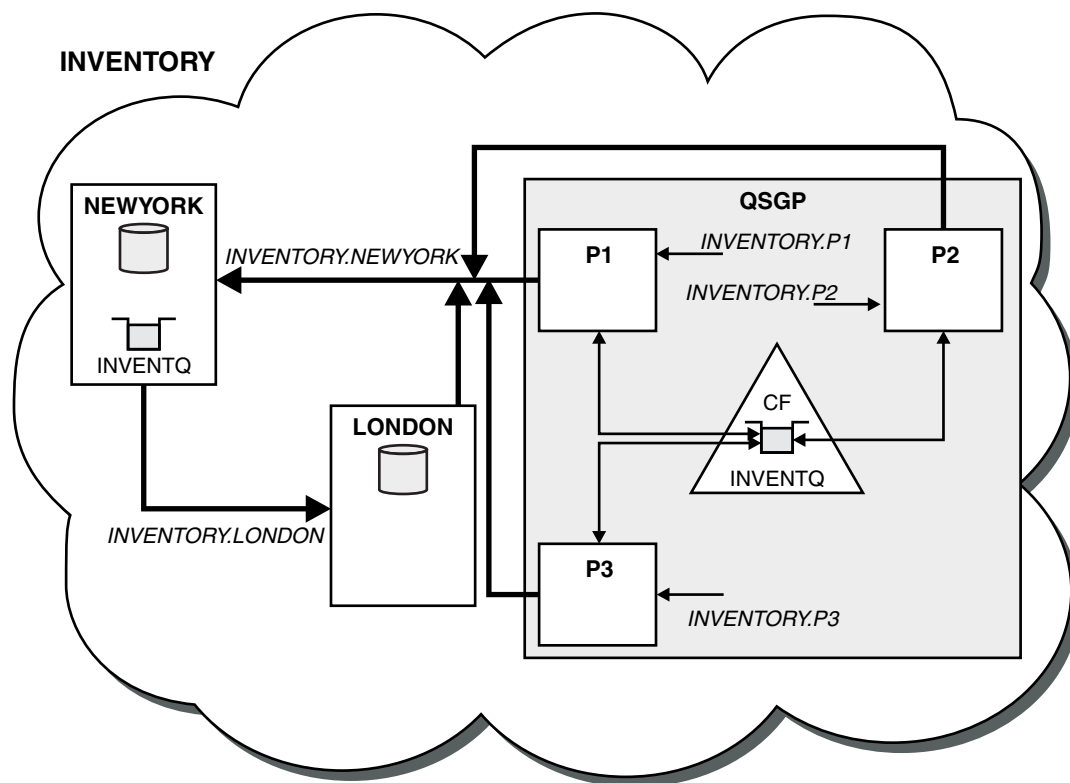


Figure 124. Cluster and queue-sharing group

Now messages put on the INVENTQ queue by LONDON are routed alternately around the four queue managers advertised as hosting the queue.

## What to do next

A benefit of having members of a queue-sharing group host a cluster queue is any member of the group can reply to a request. In this case perhaps P1 becomes unavailable after receiving a message on the shared queue. Another member of the queue-sharing group can reply instead.

## Moving a full repository to another queue manager

Move a full repository from one queue manager to another, building up the new repository from information held at the second repository.

## Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager to a cluster” on page 924.
- For business reasons you now want to remove the full repository from queue manager LONDON, and replace it with a full repository at queue manager PARIS. The NEWYORK queue manager is to continue holding a full repository.

## About this task

Follow these steps to move a full repository to another queue manager.

## Procedure

1. Alter PARIS to make it a full repository queue manager.

On PARIS, issue the following command:

```
ALTER QMGR REPOS(INVENTORY)
```

2. Add a CLUSSDR channel on PARIS

PARIS currently has a cluster-sender channel pointing to LONDON. LONDON is no longer to hold a full repository for the cluster. PARIS must have a new cluster-sender channel that points to NEWYORK, where the other full repository is now held.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from PARIS to repository at NEWYORK')
```

3. Define a CLUSSDR channel on NEWYORK that points to PARIS

Currently NEWYORK has a cluster-sender channel pointing to LONDON. Now that the other full repository has moved to PARIS, you need to add a new cluster-sender channel at NEWYORK that points to PARIS.

```
DEFINE CHANNEL(INVENTORY.PARIS) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(PARIS.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from NEWYORK to repository at PARIS')
```

When you add the cluster-sender channel to PARIS, PARIS learns about the cluster from NEWYORK. It builds up its own full repository using the information from NEWYORK.

4. Check that queue manager PARIS now has a full repository

Check that queue manager PARIS has built its own full repository from the full repository on queue manager NEWYORK. Issue the following commands:

```
DIS QCLUSTER(*) CLUSTER (INVENTORY)
DIS CLUSQMGR(*) CLUSTER (INVENTORY)
```

Check that these commands show details of the same resources in this cluster as on NEWYORK.

**Note:** If queue manager NEWYORK is not available, this building of information cannot complete. Do not move on to the next step until the task is complete.

5. Alter the queue-manager definition on LONDON

Finally alter the queue manager at LONDON so that it no longer holds a full repository for the cluster. On LONDON, issue the command:

```
ALTER QMGR REPOS(' ')
```

The queue manager no longer receives any cluster information. After 30 days the information that is stored in its full repository expires. The queue manager LONDON now builds up its own partial repository.

6. Remove or change any outstanding definitions.

When you are sure that the new arrangement of your cluster is working as expected, remove or change manually defined CLUSSDR definitions that are no longer correct.

- On the PARIS queue manager, you must stop and delete the cluster-sender channel to LONDON, and then issue the start channel command so that the cluster can use the automatic channels again:

```
STOP CHANNEL(INVENTORY.LONDON)
DELETE CHANNEL(INVENTORY.LONDON)
START CHANNEL(INVENTORY.LONDON)
```

- On the NEWYORK queue manager, you must stop and delete the cluster-sender channel to LONDON, and then issue the start channel command so that the cluster can use the automatic channels again:

```
STOP CHANNEL(INVENTORY.LONDON)
DELETE CHANNEL(INVENTORY.LONDON)
START CHANNEL(INVENTORY.LONDON)
```

- Replace all other manually defined cluster-sender channels that point to LONDON on all queue managers in the cluster with channels that point to either NEWYORK or PARIS. After deleting a channel, always issue the **start channel** command so that the cluster can use the automatic channels again. In this small example, there are no others. To check whether there are any others that you have forgotten, issue the DISPLAY CHANNEL command from each queue manager, specifying TYPE(CLUSSDR). For example:

```
DISPLAY CHANNEL(*) TYPE(CLUSSDR)
```

It is important that you perform this task as soon as possible after moving the full repository from LONDON to PARIS. In the time before you perform this task, queue managers that have manually defined CLUSSDR channels named INVENTORY.LONDON might send requests for information using this channel.

After LONDON has ceased to be a full repository, if it receives such requests it will write error messages to its queue manager error log. The following examples show which error messages might be seen on LONDON:

- AMQ9428: Unexpected publication of a cluster queue object received
- AMQ9432: Query received by a non-repository queue manager

The queue manager LONDON does not respond to the requests for information because it is no longer a full repository. The queue managers requesting information from LONDON must rely on NEWYORK for cluster information until their manually defined CLUSSDR definitions are corrected to point to PARIS. This situation must not be tolerated as a valid configuration in the long term.

## Results

Figure 125 on page 937 shows the cluster set up by this task.



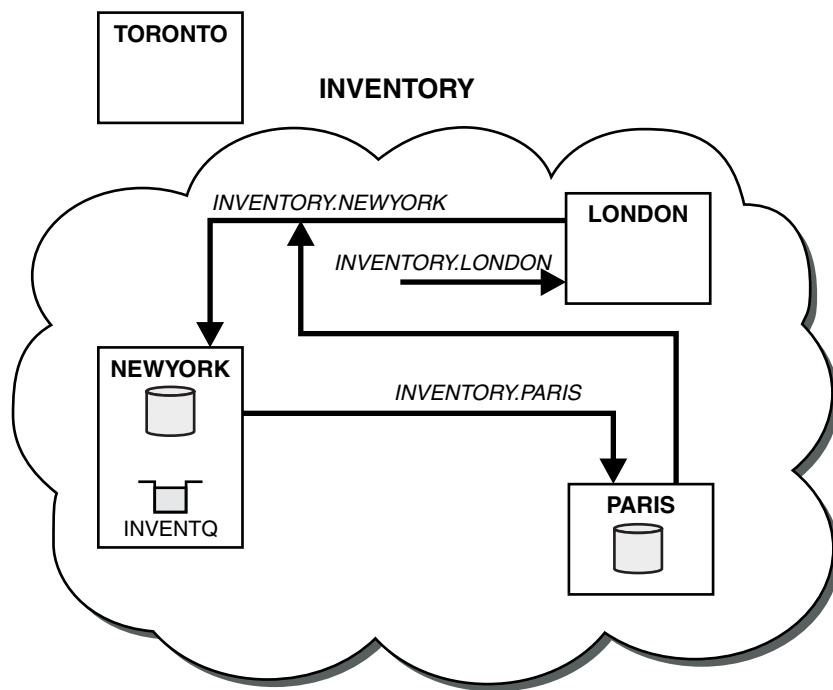


Figure 125. The *INVENTORY* cluster with the full repository moved to *PARIS*

## Establishing communication in a cluster

A channel initiator is needed to start a communication channel when there is a message to deliver. A channel listener waits to start the other end of a channel to receive the message.

### Before you begin

To establish communication between queue managers in a cluster, configure a link using one of the supported communication protocols. The supported protocols are TCP or LU 6.2 on any platform, and NetBIOS or SPX on Windows systems. As part of this configuration, you also need channel initiators and channel listeners just as you do with distributed queuing.

### About this task

All cluster queue managers need a channel initiator to monitor the system-defined initiation queue `SYSTEM.CHANNEL.INITQ`. `SYSTEM.CHANNEL.INITQ` is the initiation queue for all transmission queues including the cluster transmission queue.

Each queue manager must have a channel listener. A channel listener program waits for incoming network requests and starts the appropriate receiver-channel when it is needed. The implementation of channel listeners is platform-specific, however there are some common features. On all IBM MQ platforms, the listener can be started using the `START LISTENER` command. On IBM MQ for IBM i, Windows, UNIX and Linux systems, you can start the listener automatically at the same time as the queue manager. To start the listener automatically, set the `CONTROL` attribute of the `LISTENER` object to `QMGR` or `STARTONLY`.

### Procedure

1. Start the channel initiator.

-  z/OS

## IBM MQ for z/OS

There is one channel initiator for each queue manager and it runs as a separate address space. You start it using the **MQSC** `START CHINIT` command, which you issue as part of your queue manager startup.

-   

## IBM MQ for Windows, UNIX and Linux systems

When you start a queue manager, if the queue manager attribute `SCHINIT` is set to `QMGR`, a channel initiator is automatically started. Otherwise it can be started using the **runmqsc** `START CHINIT` command or the **runmqchi** control command.

- 

## IBM MQ for IBM i

When you start a queue manager, if the queue manager attribute `SCHINIT` is set to `QMGR`, a channel initiator is automatically started. Otherwise it can be started using the **runmqsc** `START CHINIT` command or the **runmqchi** control command.

## 2. Start the channel listener.

- 

## IBM MQ for z/OS

Use the channel listener program provided by IBM MQ. To start an IBM MQ channel listener, use the **MQSC** command `START LISTENER`, which you issue as part of your channel initiator startup. For example:

```
START LISTENER PORT(1414) TRPTYPE(TCP)
```

or:

```
START LISTENER LUNAME(LONDON.LUNAME) TRPTYPE(LU62)
```

Members of a queue-sharing group can use a shared listener instead of a listener for each queue manager. Do not use shared listeners with clusters. Specifically, do not make the `CONNNAME` of the `CLUSRCVR` channel the address of the shared listener of the queue sharing group. If you do, queue managers might receive messages for queues for which they do not have a definition.

- 

## IBM MQ for IBM i

Use the channel listener program provided by IBM MQ. To start an IBM MQ channel listener use the **CL** command `STRMQMLSR`. For example:

```
STRMQMLSR MQMNAME(QM1) PORT(1414)
```

- 

## IBM MQ for Windows

Use either the channel listener program provided by IBM MQ, or the facilities provided by the operating system.

To start the IBM MQ channel listener use the `RUNMQLSR` command. For example:

```
RUNMQLSR -t tcp -p 1414 -m QM1
```

-  

## IBM MQ on UNIX and Linux systems

Use either the channel listener program provided by IBM MQ, or the facilities provided by the operating system; for example, **inetd** for TCP communications.

To start the IBM MQ channel listener use the **runmq1sr** command. For example:

```
runmq1sr -t tcp -p 1414 -m QM1
```

To use **inetd** to start channels, configure two files:

- a. Edit the file `/etc/services`. You must be logged in as a superuser or root. If the following line is not in the file, add it as shown:

```
MQSeries 1414/tcp # Websphere MQ channel listener
```

where 1414 is the port number required by IBM MQ. You can change the port number, but it must match the port number specified at the sending end.

- b. Edit the file `/etc/inetd.conf`. If you do not have the following line in that file, add it as shown:

```
MQSeries stream tcp nowait mqm MQ_INSTALLATION_PATH/bin/amqcrsta amqcrsta
-m queue.manager.name
```

where `MQ_INSTALLATION_PATH` is replaced by the high-level directory in which IBM MQ is installed.

The updates become active after **inetd** has reread the configuration files. Issue the following commands from the root user ID:

On AIX:

```
refresh -s inetd
```

On HP-UX:

```
inetd -c
```

On Solaris or Linux:

- a. Find the process ID of the **inetd** with the command:

```
ps -ef | grep inetd
```

- b. Run the appropriate command, as follows:

– For Solaris 9 and Linux:

```
kill -1 inetd processid
```

– For Solaris 10, or later versions:

```
inetconv
```

## Converting an existing network into a cluster

Convert an existed distributed queuing network to a cluster and add an additional queue manager to increase capacity.

### Before you begin

In “Setting up a new cluster” on page 913 through “Moving a full repository to another queue manager” on page 935 you created and extended a new cluster. The next two tasks explore a different approach: that of converting an existing network of queue managers into a cluster.

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- A IBM MQ network is already in place, connecting the nationwide branches of a chain store. It has a hub and spoke structure: all the queue managers are connected to one central queue manager. The central queue manager is on the system on which the inventory application runs. The application is driven by the arrival of messages on the `INVENTQ` queue, for which each queue manager has a remote-queue definition.

This network is illustrated in Figure 126 on page 940.




Figure 126. A hub and spoke network

- To ease administration you are going to convert this network into a cluster and create another queue manager at the central site to share the workload.

The cluster name is CHNSTORE.

**Note:** The cluster name CHNSTORE was selected to allow cluster-receiver channel names to be created using names in the format *cluster-name.queue-manager* that do not exceed the maximum length of 20 characters, for example CHNSTORE.WASHINGTON.

- Both the central queue managers are to host full repositories and be accessible to the inventory application.
- The inventory application is to be driven by the arrival of messages on the INVENTQ queue hosted by either of the central queue managers.
- The inventory application is to be the only application running in parallel and accessible by more than one queue manager. All other applications continue to run as before.
- All the branches have network connectivity to the two central queue managers.
- The network protocol is TCP.

## About this task

Follow these steps to convert an existing network into a cluster.

### Procedure

1. Review the inventory application for message affinities.

Before proceeding ensure that the application can handle message affinities. Message affinities are the relationship between conversational messages that are exchanged between two applications, where the messages must be processed by a particular queue manager or in a particular sequence. For more information on message affinities, see: "Handling message affinities" on page 1001

2. Alter the two central queue managers to make them full repository queue managers.

The two queue managers CHICAGO and CHICAGO2 are at the hub of this network. You have decided to concentrate all activity associated with the chain store cluster on to those two queue managers. As well as the inventory application and the definitions for the INVENTQ queue, you want these queue managers to host the two full repositories for the cluster. At each of the two queue managers, issue the following command:

```
ALTER QMGR REPOS(CHNSTORE)
```

3. Define a CLUSRCVR channel on each queue manager.

At each queue manager in the cluster, define a cluster-receiver channel and a cluster-sender channel. It does not matter which channel you define first.

Make a CLUSRCVR definition to advertise each queue manager, its network address, and other information, to the cluster. For example, on queue manager ATLANTA:

```
DEFINE CHANNEL(CHNSTORE.ATLANTA) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(ATLANTA.CHSTORE.COM) CLUSTER(CHNSTORE)
DESCR('Cluster-receiver channel')
```

4. Define a CLUSSDR channel on each queue manager

Make a CLUSSDR definition at each queue manager to link that queue manager to one or other of the full repository queue managers. For example, you might link ATLANTA to CHICAGO2:

```
DEFINE CHANNEL(CHNSTORE.CHICAGO2) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(CHICAGO2.CHSTORE.COM) CLUSTER(CHNSTORE)
DESCR('Cluster-sender channel to repository queue manager')
```

5. Install the inventory application on CHICAGO2.

You already have the inventory application on queue manager CHICAGO. Now you need to make a copy of this application on queue manager CHICAGO2.

6. Define the INVENTQ queue on the central queue managers.

On CHICAGO, modify the local queue definition for the queue INVENTQ to make the queue available to the cluster. Issue the command:

```
ALTER QLOCAL(INVENTQ) CLUSTER(CHNSTORE)
```

On CHICAGO2, make a definition for the same queue:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(CHNSTORE)
```

On z/OS, you can use the MAKEDEF option of the COMMAND function of **CSQUTIL** to make an exact copy on CHICAGO2 of the INVENTQ on CHICAGO.

When you make these definitions, a message is sent to the full repositories at CHICAGO and CHICAGO2 and the information in them is updated. The queue manager finds out from the full repositories when it puts a message to the INVENTQ, that there is a choice of destinations for the messages.

7. Check that the cluster changes have been propagated.

Check that the definitions you created in the previous step have been propagated through the cluster. Issue the following command on a full repository queue manager:

```
DIS QCLUSTER(INVENTQ)
```

## Adding a new, interconnected cluster:

Add a new cluster that shares some queue managers with an existing cluster.

### Before you begin

#### Note:

1. For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.
2. Before starting this task, check for queue-name clashes and understand the consequences. You might need to rename a queue, or set up queue aliases before you can proceed.

#### Scenario:

- An IBM MQ cluster has been set up as described in “Converting an existing network into a cluster” on page 939.
- A new cluster called MAILORDER is to be implemented. This cluster comprises four of the queue managers that are in the CHNSTORE cluster; CHICAGO, CHICAGO2, SEATTLE, and ATLANTA, and two additional queue managers; HARTFORD and OMAHA. The MAILORDER application runs on the system at Omaha, connected to queue manager OMAHA. It is driven by the other queue managers in the cluster putting messages on the MORDERQ queue.
- The full repositories for the MAILORDER cluster are maintained on the two queue managers CHICAGO and CHICAGO2.
- The network protocol is TCP.

### About this task

Follow these steps to add a new, interconnected cluster.

### Procedure

1. Create a namelist of the cluster names.

The full repository queue managers at CHICAGO and CHICAGO2 are now going to hold the full repositories for both of the clusters CHNSTORE and MAILORDER. First, create a namelist containing the names of the clusters. Define the namelist on CHICAGO and CHICAGO2, as follows:

```
DEFINE NAMLIST(CHAINMAIL)
DESCR('List of cluster names')
NAMES(CHNSTORE, MAILORDER)
```

2. Alter the two queue-manager definitions.

Now alter the two queue-manager definitions at CHICAGO and CHICAGO2. Currently these definitions show that the queue managers hold full repositories for the cluster CHNSTORE. Change that definition to show that the queue managers hold full repositories for all clusters listed in the CHAINMAIL namelist. Alter the CHICAGO and CHICAGO2 queue manager definitions:

```
ALTER QMGR REPOS(' ') REPOSNL(CHAINMAIL)
```

3. Alter the CLUSRCVR channels on CHICAGO and CHICAGO2.

The CLUSRCVR channel definitions at CHICAGO and CHICAGO2 show that the channels are available in the cluster CHNSTORE. You need to change the cluster-receiver definition to show that the channels are available to all clusters listed in the CHAINMAIL namelist. Change the cluster-receiver definition at CHICAGO:

```
ALTER CHANNEL(CHNSTORE.CHICAGO) CHLTYPE(CLUSRCVR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

At CHICAGO2, enter the command:

```
ALTER CHANNEL(CHNSTORE.CHICAGO2) CHLTYPE(CLUSRCVR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

4. Alter the CLUSSDR channels on CHICAGO and CHICAGO2.

Change the two CLUSSDR channel definitions to add the namelist. At CHICAGO, enter the command:

```
ALTER CHANNEL(CHNSTORE.CHICAGO2) CHLTYPE(CLUSSDR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

At CHICAGO2, enter the command:

```
ALTER CHANNEL(CHNSTORE.CHICAGO) CHLTYPE(CLUSSDR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

5. Create a namelist on SEATTLE and ATLANTA.

Because SEATTLE and ATLANTA are going to be members of more than one cluster, you must create a namelist containing the names of the clusters. Define the namelist on SEATTLE and ATLANTA, as follows:

```
DEFINE NAMELIST(CHAINMAIL)
DESCR('List of cluster names')
NAMES(CHNSTORE, MAILORDER)
```

6. Alter the CLUSRCVR channels on SEATTLE and ATLANTA.

The CLUSRCVR channel definitions at SEATTLE and ATLANTA show that the channels are available in the cluster CHNSTORE. Change the cluster-receive channel definitions to show that the channels are available to all clusters listed in the CHAINMAIL namelist. At SEATTLE, enter the command:

```
ALTER CHANNEL(CHNSTORE.SEATTLE) CHLTYPE(CLUSRCVR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

At ATLANTA, enter the command:

```
ALTER CHANNEL(CHNSTORE.ATLANTA) CHLTYPE(CLUSRCVR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

7. Alter the CLUSSDR channels on SEATTLE and ATLANTA.

Change the two CLUSSDR channel definitions to add the namelist. At SEATTLE, enter the command:

```
ALTER CHANNEL(CHNSTORE.CHICAGO) CHLTYPE(CLUSSDR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

At ATLANTA, enter the command:

```
ALTER CHANNEL(CHNSTORE.CHICAGO2) CHLTYPE(CLUSSDR)
CLUSTER(' ') CLUSNL(CHAINMAIL)
```

8. Define CLUSRCVR and CLUSSDR channels on HARTFORD and OMAHA.

On the two new queue managers HARTFORD and OMAHA, define cluster-receiver and cluster-sender channels. It does not matter in which sequence you make the definitions. At HARTFORD, enter:

```
DEFINE CHANNEL(MAILORDER.HARTFORD) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(HARTFORD.CHSTORE.COM) CLUSTER(MAILORDER)
DESCR('Cluster-receiver channel for HARTFORD')
```

```
DEFINE CHANNEL(MAILORDER.CHICAGO) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(CHICAGO.CHSTORE.COM) CLUSTER(MAILORDER)
DESCR('Cluster-sender channel from HARTFORD to repository at CHICAGO')
```

At OMAHA, enter:

```
DEFINE CHANNEL(MAILORDER.OMAHA) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(OMAHA.CHSTORE.COM) CLUSTER(MAILORDER)
DESCR('Cluster-receiver channel for OMAHA')
```

```
DEFINE CHANNEL(MAILORDER.CHICAGO) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(CHICAGO.CHSTORE.COM) CLUSTER(MAILORDER)
DESCR('Cluster-sender channel from OMAHA to repository at CHICAGO')
```

9. Define the MORDERQ queue on OMAHA.

The final step to complete this task is to define the queue MORDERQ on the queue manager OMAHA. At OMAHA, enter:

```
DEFINE QLOCAL(MORDERQ) CLUSTER(MAILORDER)
```

10. Check that the cluster changes have been propagated.

Check that the definitions you created with the previous steps have been propagated through the cluster. Issue the following commands on a full repository queue manager:

```
DIS QCLUSTER (MORDERQ)
DIS CLUSQMGR
```

11.

## Results

The cluster set up by this task is shown in Figure 127 on page 945.

Now we have two overlapping clusters. The full repositories for both clusters are held at CHICAGO and CHICAGO2. The mail order application that runs on OMAHA is independent of the inventory application that runs at CHICAGO. However, some of the queue managers that are in the CHNSTORE cluster are also in the MAILORDER cluster, and so they can send messages to either application. Before carrying out this task to overlap two clusters, be aware of the possibility of queue-name clashes.

Suppose that on NEWYORK in cluster CHNSTORE and on OMAHA in cluster MAILORDER, there is a queue called ACCOUNTQ. If you overlap the clusters and then an application on SEATTLE puts a message to the queue ACCOUNTQ, the message can go to either instance of the ACCOUNTQ.



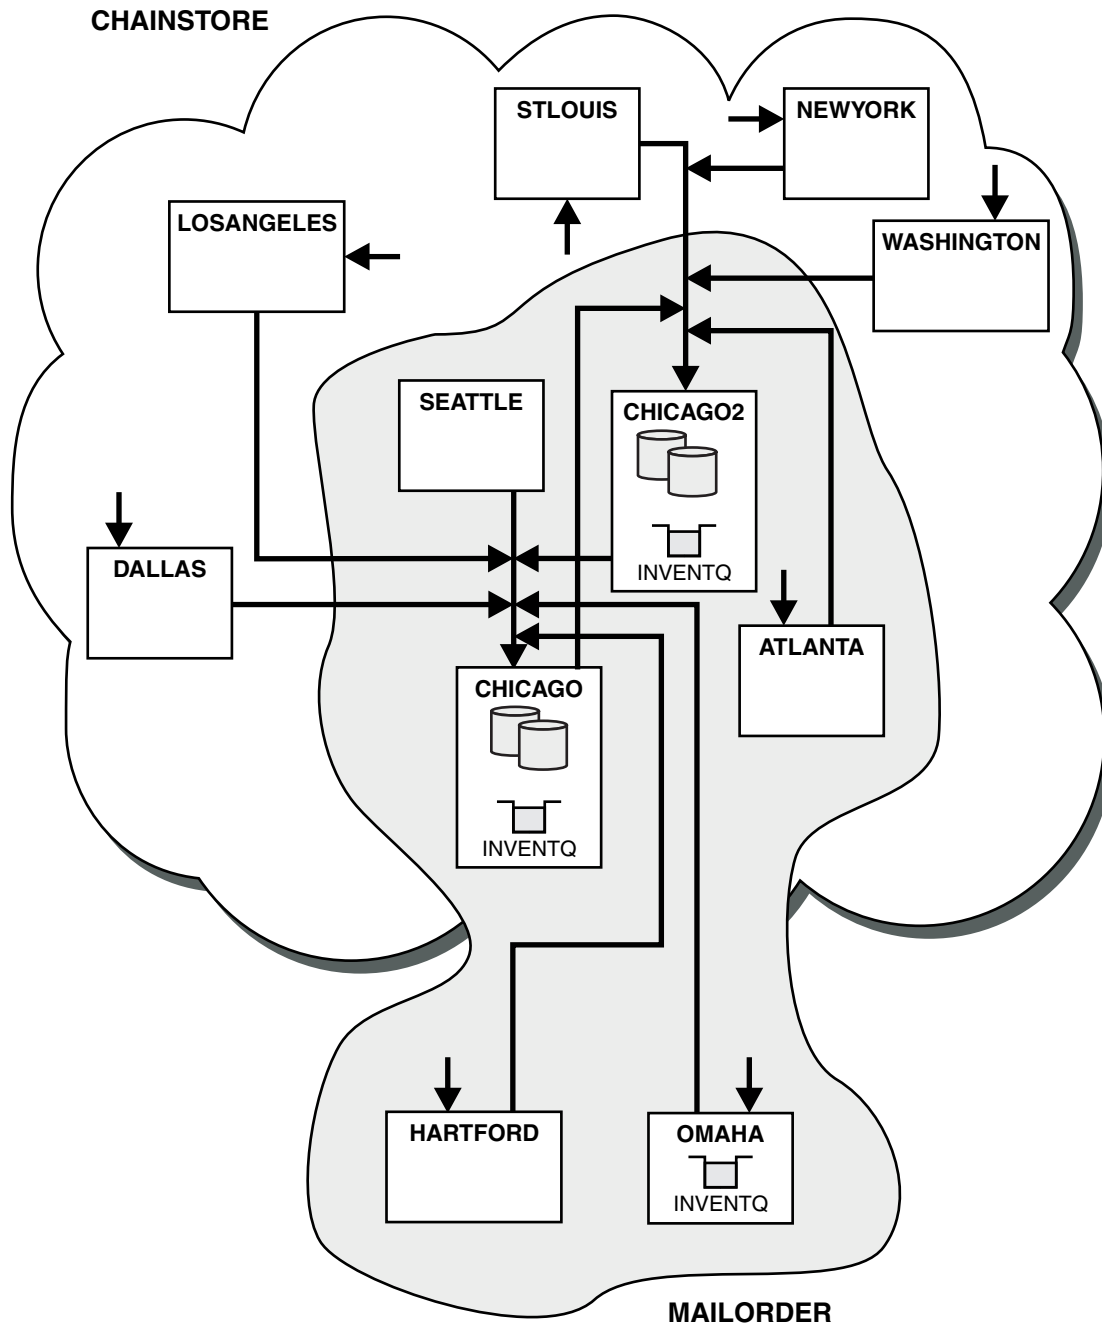


Figure 127. Interconnected clusters

### What to do next

Suppose you decide to merge the MAILORDER cluster with the CHNSTORE cluster to form one large cluster called CHNSTORE.

To merge the MAILORDER cluster with the CHNSTORE cluster, such that CHICAGO and CHICAGO2 hold the full repositories:

- Alter the queue manager definitions for CHICAGO and CHICAGO2, removing the REPOSNL attribute, which specifies the namelist ( CHAINMAIL), and replacing it with a REPOS attribute specifying the cluster name ( CHNSTORE). For example:

```
ALTER QMGR(CHICAGO) REPOSNL(' ') REPOS(CHNSTORE)
```

- On each queue manager in the MAILORDER cluster, alter all the channel definitions and queue definitions to change the value of the CLUSTER attribute from MAILORDER to CHNSTORE. For example, at HARTFORD, enter:

```
ALTER CHANNEL(MAILORDER.HARTFORD) CLUSTER(CHNSTORE)
```

At OMAHA enter:

```
ALTER QLOCAL(MORDERQ) CLUSTER(CHNSTORE)
```

- Alter all definitions that specify the cluster namelist CHAINMAIL, that is, the CLUSRCVR and CLUSSDR channel definitions at CHICAGO, CHICAGO2, SEATTLE, and ATLANTA, to specify instead the cluster CHNSTORE.

From this example, you can see the advantage of using namelists. Instead of altering the queue manager definitions for CHICAGO and CHICAGO2 you can alter the value of the namelist CHAINMAIL. Similarly, instead of altering the CLUSRCVR and CLUSSDR channel definitions at CHICAGO, CHICAGO2, SEATTLE, and ATLANTA, you can achieve the required result by altering the namelist.

### Removing a cluster network:

Remove a cluster from a network and restore the distributed queuing configuration.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- A IBM MQ cluster has been set up as described in “Converting an existing network into a cluster” on page 939.
- This cluster is now to be removed from the system. The network of queue managers is to continue functioning as it did before the cluster was implemented.

### About this task

Follow these steps to remove a cluster network.

### Procedure

1. Remove cluster queues from the CHNSTORE cluster.

On both CHICAGO and CHICAGO2, modify the local queue definition for the queue INVENTQ to remove the queue from the cluster. Issue the command:

```
ALTER QLOCAL(INVENTQ) CLUSTER(' ')
```

When you alter the queue, the information in the full repositories is updated and propagated throughout the cluster. Active applications using MQ00\_BIND\_NOT\_FIXED, and applications using MQ00\_BIND\_AS\_Q\_DEF where the queue has been defined with DEFBIND(NOTFIXED), fail on the next attempted MQPUT or MQPUT1 call. The reason code MQRC\_UNKNOWN\_OBJECT\_NAME is returned.

You do not have to perform Step 1 first, but if you do not, perform it instead after Step 4.

2. Stop all applications that have access to cluster queue.

Stop all applications that have access to cluster queues. If you do not, some cluster information might remain on the local queue manager when you refresh the cluster in Step 5. This information is removed when all applications have stopped and the cluster channels have disconnected.

3. Remove the repository attribute from the full repository queue managers.

On both CHICAGO and CHICAGO2, modify the queue manager definitions to remove the repository attribute. To do this issue the command:

```
ALTER QMGR REPOS(' ')
```

The queue managers inform the other queue managers in the cluster that they no longer hold the full repositories. When the other queue managers receive this information, you see a message indicating that the full repository has ended. You also see one or more messages indicating that there are no longer any repositories available for the cluster CHNSTORE.

4. Remove cluster channels.

On CHICAGO remove the cluster channels:

```
ALTER CHANNEL(CHNSTORE.CHICAGO2) CHLTYPE(CLUSSDR) CLUSTER(' ')
ALTER CHANNEL(CHNSTORE.CHICAGO) CHLTYPE(CLUSRCVR) CLUSTER(' ')
```

**Note:** It is important to issue the CLUSSDR command first, then CLUSRCVR command. Do not issue the CLUSRCVR command first, then the CLUSSDR command. Doing so, creates indoubt channels that have a STOPPED status. You then need to issue a START CHANNEL command to recover the stopped channels; for example, START CHANNEL(CHNSTORE.CHICAGO).

You see messages indicating that there are no repositories for the cluster CHNSTORE.

If you did not remove the cluster queues as described in Step 1, do so now.

5. Stop cluster channels.

On CHICAGO stop the cluster channels with the following commands:

```
STOP CHANNEL(CHNSTORE.CHICAGO2)
STOP CHANNEL(CHNSTORE.CHICAGO)
```

6. Repeat steps 4 and 5 for each queue manager in the cluster.
7. Stop the cluster channels, then remove all definitions for the cluster channels and cluster queues from each queue manager.
8. Optional: Clear the cached cluster information held by the queue manager. Although the queue managers are no longer members of the cluster, they each retain a cached copy of information about the cluster. If you want to remove this data, see task “Restoring a queue manager to its pre-cluster state” on page 971.
9. Replace the remote-queue definitions for the INVENTQ  
So that the network can continue to function, replace the remote queue definition for the INVENTQ at every queue manager.
10. Tidy up the cluster.  
Delete any queue or channel definitions no longer required.

## Creating two-overlapping clusters with a gateway queue manager

Follow the instructions in the task to construct overlapping clusters with a gateway queue manager. Use the clusters as a starting point for the following examples of isolating messages to one application from messages to other applications in a cluster.

### About this task

The example cluster configuration used to illustrate isolating cluster message traffic is shown in Figure 128 on page 948. The example is described in Clustering: Application isolation using multiple cluster transmission queues.

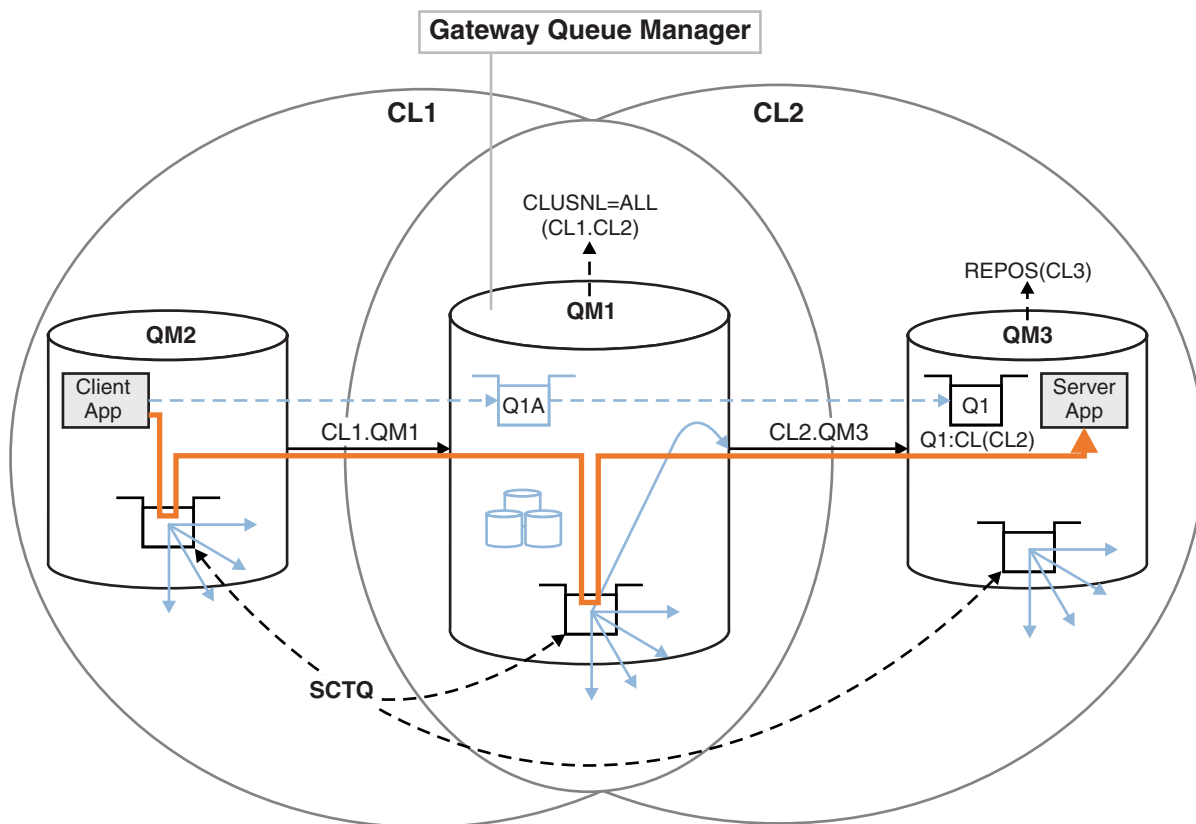


Figure 128. Client-server application deployed to hub and spoke architecture using IBM MQ clusters

To make the number of steps to construct the example as few as possible, the configuration is kept simple, rather than realistic. The example might represent the integration of two clusters created by two separate organizations. For a more realistic scenario, see Clustering: Planning how to configure cluster transmission queues.

Follow the steps to construct the clusters. The clusters are used in the following examples of isolating the message traffic from the client application to the server application.

The instructions add a couple of extra queue managers so that each cluster has two repositories. The gateway queue manager is not used as a repository for performance reasons.

## Procedure

1. Create and start the queue managers QM1, QM2, QM3, QM4, QM5.

```
crtmqm -sax -u SYSTEM.DEAD.LETTER.QUEUE QM n
strmqm QmgrName
```

**Note:** QM4 and QM5 are the backup full repositories for the clusters.

2. Define and start listeners for each of the queue managers.

```
*... On QM n
DEFINE LISTENER(TCP141 n) TRPTYPE(TCP) IPADDR(hostname) PORT(141 n) CONTROL(QMGR) REPLACE
START LISTENER(TCP141 n)
```

3. Create a cluster name list for all of the clusters.

```
*... On QM1
DEFINE NAMELIST(ALL) NAMES(CL1, CL2) REPLACE
```

4. Make QM2 and QM4 full repositories for CL1, QM3 and QM5 full repositories for CL2.

- a. For CL1:
  - \*... On QM2 and QM4
  - ALTER QMGR REPOS(CL1) DEFCLXQ(SCTQ)
- b. For CL2:
  - \*... On QM3 and QM5
  - ALTER QMGR REPOS(CL2) DEFCLXQ(SCTQ)

5. Add the cluster-sender and cluster-receiver channels for each queue manager and cluster.  
Run the following commands on QM2, QM3, QM4 and QM5, where *c*, *n*, and *m* take the values shown in Table 134 for each queue manager:

Table 134. Parameter values for creating clusters 1 and 2

| Queue manager | Cluster<br><i>c</i> | Other repository<br><i>n</i> | This repository<br><i>m</i> |
|---------------|---------------------|------------------------------|-----------------------------|
| QM2           | 1                   | 4                            | 2                           |
| QM4           | 1                   | 2                            | 4                           |
| QM3           | 2                   | 5                            | 3                           |
| QM5           | 2                   | 3                            | 5                           |

```
*... On QM m
DEFINE CHANNEL(CL c.QM n) CHLTYPE(CLUSSDR) CONNAME('localhost(141 n)') CLUSTER(CL c) REPLACE
DEFINE CHANNEL(CL c.QM m) CHLTYPE(CLUSRCVR) CONNAME('localhost(141 m)') CLUSTER(CL c) REPLACE
```

6. Add the gateway queue manager, QM1, to each of the clusters.

```
*... On QM1
DEFINE CHANNEL(CL1.QM2) CHLTYPE(CLUSSDR) CONNAME('localhost(1412)') CLUSTER(CL1) REPLACE
DEFINE CHANNEL(CL1.QM1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1411)') CLUSTER(CL1) REPLACE
DEFINE CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR) CONNAME('localhost(1413)') CLUSTER(CL2) REPLACE
DEFINE CHANNEL(CL2.QM1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1411)') CLUSTER(CL2) REPLACE
```

7. Add the local queue Q1 to queue manager QM3 in cluster CL2.

```
*... On QM3
DEFINE QLOCAL(Q1) CLUSTER(CL2) REPLACE
```

8. Add the clustered queue manager alias Q1A to the gateway queue manager.

```
*... On QM1
DEFINE QALIAS(Q1A) CLUSNL(ALL) TARGET(Q1) TARGTYPE(Queue) DEFBIND(NOTFIXED) REPLACE
```

**Note:** Applications using the queue manager alias on any other queue manager but QM1, must specify DEFBIND(NOTFIXED) when they open the alias queue. **DEFBIND** specifies whether the routing information in the message header is fixed when the queue is opened by the application. If it is set to the default value, OPEN, messages are routed to Q1@QM1. Q1@QM1 does not exist, so messages from other queue managers end up on a dead letter queue. By setting the queue attribute to DEFBIND(NOTFIXED), applications such as **amqsput**, which default to the queue setting of **DEFBIND**, behave in the correct way.

9. Add the cluster queue manager alias definitions for all the clustered queue managers to the gateway queue manager, QM1.

```
*... On QM1
DEFINE QREMOTE(QM2) RNAME(' ') RQMNAME(QM2) CLUSNL(ALL) REPLACE
DEFINE QREMOTE(QM3) RNAME(' ') RQMNAME(QM3) CLUSNL(ALL) REPLACE
```

**Tip:** The queue manager alias definitions on the gateway queue manager transfer messages that refer to a queue manager in another cluster; see Clustered queue manager aliases.

## What to do next

1. Test the queue alias definition by sending a message from QM2 to Q1 on QM3 using the queue alias definition Q1A.
  - a. Run the sample program **amqsput** on QM2 to put a message.

```

C:\IBM\MQ>amqsput Q1A QM2
Sample AMQSPUT0 start
target queue is Q1A
Sample request message from QM2 to Q1 using Q1A

```

```

Sample AMQSPUT0 end

```

- b. Run the sample program **amqsget** to get the message from Q1 on QM3

```

C:\IBM\MQ>amqsget Q1 QM3
Sample AMQSGET0 start
message <Sample request message from QM2 to Q1 using Q1A>
no more messages
Sample AMQSGET0 end

```

2. Test the queue manager alias definitions by sending a request message and receiving a reply message on a temporary-dynamic reply queue.

The diagram shows the path taken by the reply message back to a temporary dynamic queue, which is called RQ. The server application, connected to QM3, opens the reply queue using the queue manager name QM2. The queue manager name QM2 is defined as a clustered queue manager alias on QM1. QM3 routes the reply message to QM1. QM1 routes the message to QM2.

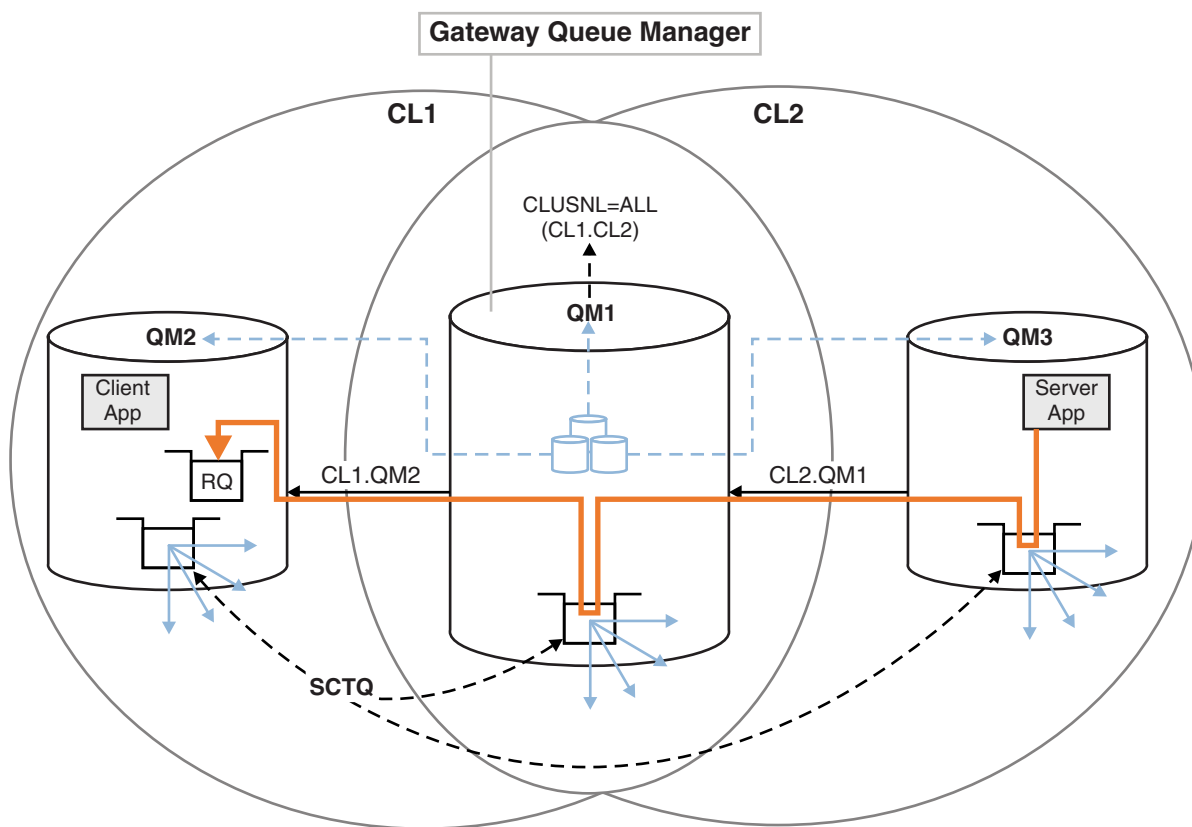


Figure 129. Using a queue manager alias to return the reply message to a different cluster

The way the routing works is as follows. Every queue manager in each cluster has a queue manager alias definition on QM1. The aliases are clustered in all the clusters. The grey dashed arrows from each of the aliases to a queue manager show that each queue manager alias is resolved to a real queue manager in at least one of the clusters. In this case, the QM2 alias is clustered in both cluster CL1 and CL2, and is resolved to the real queue manager QM2 in CL1. The server application creates the reply message using the reply to queue name RQ, and reply to queue manager name QM2. The message is

routed to QM1 because the queue manager alias definition QM2 is defined on QM1 in cluster CL2 and queue manager QM2 is not in cluster CL2. As the message cannot be sent to the target queue manager, it is sent to the queue manager that has the alias definition.

QM1 places the message on the cluster transmission queue on QM1 for transferal to QM2. QM1 routes the message to QM2 because the queue manager alias definition on QM1 for QM2 defines QM2 as the real target queue manager. The definition is not circular, because alias definitions can refer only to real definitions; the alias cannot point to itself. The real definition is resolved by QM1, because both QM1 and QM2 are in the same cluster, CL1. QM1 finds out the connection information for QM2 from the repository for CL1, and routes the message to QM2. For the message to be rerouted by QM1, the server application must have opened the reply queue with the option DEFBIND set to MQBND\_BIND\_NOT\_FIXED. If the server application had opened the reply queue with the option MQBND\_BIND\_ON\_OPEN, the message is not rerouted and ends up on a dead letter queue.

- a. Create a clustered request queue with a trigger on QM3.

```
*... On QM3
DEFINE QLOCAL(QR) CLUSTER(CL2) TRIGGER INITQ(SYSTEM.DEFAULT.INITIATION.QUEUE) PROCESS(ECHO) REPLACE
```

- b. Create a clustered queue alias definition of QR on the gateway queue manager, QM1.

```
*... On QM1
DEFINE QALIAS(QRA) CLUSNL(ALL) TARGET(QR) TARGTYPE(QUEUE) DEFBIND(NOTFIXED) REPLACE
```

- c. Create a process definition to start the sample echo program **amqsech** on QM3.

```
*... On QM3
DEFINE PROCESS(ECHO) APPLICID(AMQSECH) REPLACE
```

- d. Create a model queue on QM2 for the sample program **amqsreq** to create the temporary-dynamic reply queue.

```
*... On QM2
DEFINE QMODEL(SYSTEM.SAMPLE.REPLY) REPLACE
```

- e. Test the queue manager alias definition by sending a request from QM2 to QR on QM3 using the queue alias definition QRA.

- 1) Run the trigger monitor program on QM3.

```
runmqtrm -m QM3
```

The output is

```
C:\IBM\MQ>runmqtrm -m QM3
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
01/02/2012 16:17:15: IBM MQ trigger monitor started.
```

---

```
01/02/2012 16:17:15: Waiting for a trigger message
```

- 2) Run the sample program **amqsreq** on QM2 to put a request and wait for a reply.

```
C:\IBM\MQ>amqsreq QRA QM2
Sample AMQSREQ0 start
server queue is QRA
replies to 4F2961C802290020
A request message from QM2 to QR on QM3
```

```
response <A request message from QM2 to QR on QM3>
no more replies
Sample AMQSREQ0 end
```

**Related tasks:**

“Adding a queue manager to a cluster: separate transmission queues” on page 926

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

**Related information:**

Access control and multiple cluster transmission queues

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

**Adding a remote queue definition to isolate messages sent from a gateway queue manager:**

Modify the configuration of overlapping clusters that use a gateway queue manager. After the modification messages are transferred to an application from the gateway queue manager without using the same transmission queue or channels as other cluster messages. The solution uses a clustered queue remote definition, and a separate sender channel and transmission queue.

**Before you begin**

Construct the overlapping clusters shown in Client-server application deployed to hub and spoke architecture using IBM MQ clusters in “Creating two-overlapping clusters with a gateway queue manager” on page 947 by following the steps in that task.

**About this task**

The solution uses distributed queuing to separate the messages for the Server App application from other message traffic on the gateway queue manager. You must define a clustered remote queue definition on QM1 to divert the messages to a different transmission queue, and a different channel. The remote queue definition must include a reference to the specific transmission queue that stores messages only for Q1 on QM3. In Figure 130 on page 953, the cluster queue alias Q1A is supplemented by a remote queue definition Q1R, and a transmission queue and sender-channel added.

In this solution, any reply messages are returned using the common `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

The advantage of this solution is that it is easy to separate traffic for multiple destination queues on the same queue manager, in the same cluster. The disadvantage of the solution is that you cannot use cluster workload balancing between multiple copies of Q1 on different queue managers. To overcome this disadvantage, see “Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954. You also have to manage the switch from one transmission queue to the other.



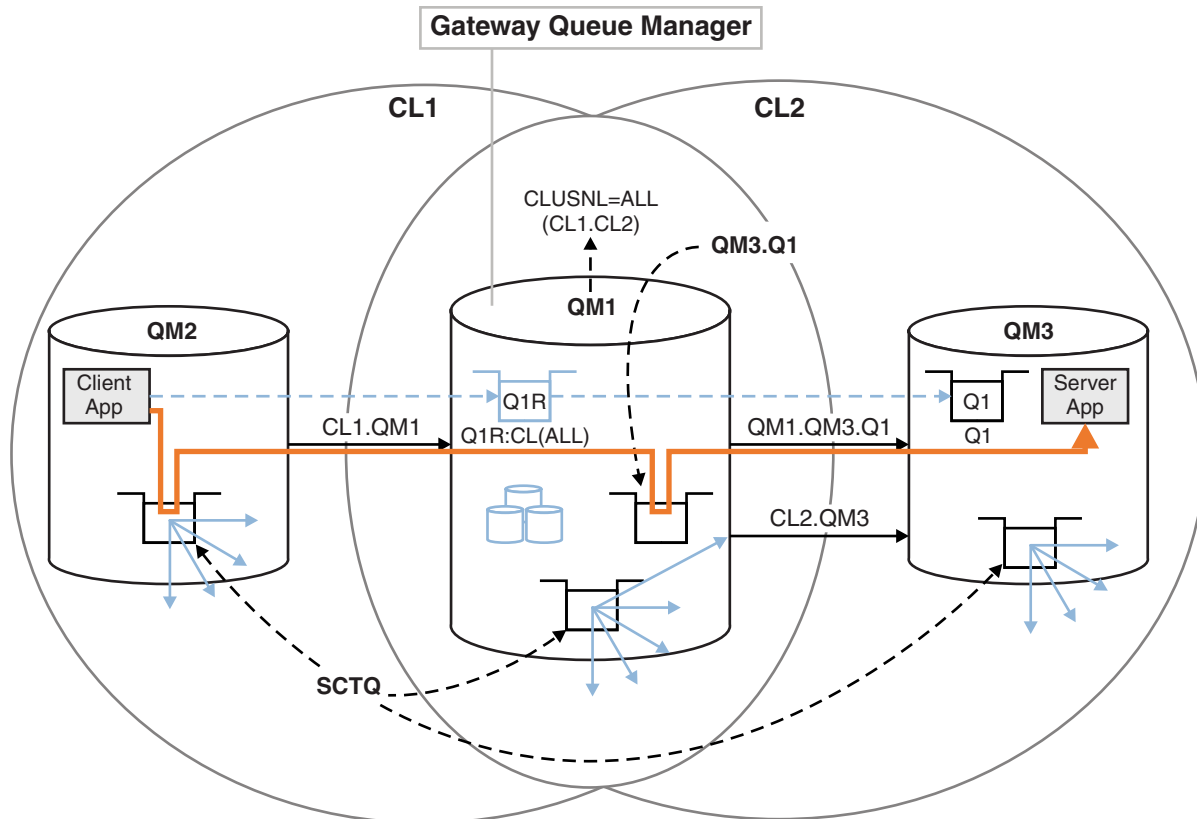


Figure 130. Client-server application deployed to hub and spoke cluster architecture using remote queue definitions

### Procedure

1. Create a channel to separate the message traffic for Q1 from the gateway queue manager
  - a. Create a sender channel on the gateway queue manager, QM1, to the target queue manager, QM3.
 

```
DEFINE CHANNEL(QM1.QM3.Q1) CHLTYPE(SDR) CONNAME(QM3HostName(1413)) XMITQ(QM3.Q1) REPLACE
```
  - b. Create a receiver channel on the target queue manager, QM3.
 

```
DEFINE CHANNEL(QM1.QM3.Q1) CHLTYPE(RCVR) REPLACE
```

2. Create a transmission queue on the gateway queue manager for message traffic to Q1
 

```
DEFINE QLOCAL(QM3.Q1) USAGE(XMITQ) REPLACE
START CHANNEL(QM1.QM3.Q1)
```

Starting the channel that is associated with the transmission queue, associates the transmission queue with the channel. The channel starts automatically, once the transmission queue has been associated with the channel.

3. Supplement the clustered queue alias definition for Q1 on the gateway queue manager with a clustered remote queue definition.
 

```
DEFINE QREMOTE CLUSNL(ALL) RNAME(Q1) RQNAME(QM3) XMITQ(QM3.Q1) REPLACE
```

### What to do next

Test the configuration by sending a message to Q1 on QM3 from QM2 using the clustered queue remote definition Q1R on the gateway queue manager QM1.

1. Run the sample program **amqsput** on QM2 to put a message.

```
C:\IBM\MQ>amqsput Q1R QM2
Sample AMQSPUT0 start
```

```
target queue is Q1R
Sample request message from QM2 to Q1 using Q1R
```

```
Sample AMQSPUT0 end
```

2. Run the sample program **amqsget** to get the message from Q1 on QM3

```
C:\IBM\MQ>amqsget Q1 QM3
Sample AMQSGET0 start
message <Sample request message from QM2 to Q1 using Q1R>
no more messages
Sample AMQSGET0 end
```

### Related tasks:

“Adding a queue manager to a cluster: separate transmission queues” on page 926

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

### Related information:

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

Access control and multiple cluster transmission queues

### Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager:

Modify the configuration of overlapping clusters that use a gateway queue manager. After the modification messages are transferred to an application from the gateway queue manager without using the same transmission queue or channels as other cluster messages. The solution uses an additional cluster transmission queue to separate message traffic to a single queue manager in a cluster.

### Before you begin

1. The gateway queue manager must be on Version 7.5, or later.
2. Construct the overlapping clusters shown in Client-server application deployed to hub and spoke architecture using IBM MQ clusters in “Creating two-overlapping clusters with a gateway queue manager” on page 947 by following the steps in that task.

### About this task

On the gateway queue manager, QM1, add a transmission queue and set its queue attribute CLCHNAME. Set CLCHNAME to the name of the cluster-receiver channel on QM3 ; see Figure 131 on page 955.

This solution has a number of advantages over the solution described in “Adding a remote queue definition to isolate messages sent from a gateway queue manager” on page 952:

- It requires fewer additional definitions.
- It supports workload balancing between multiple copies of the target queue, Q1, on different queue managers in the same cluster, CL2.
- The gateway queue manager switches automatically to the new configuration when the channel restarts without losing any messages.
- The gateway queue manager continues to forward messages in the same order as it received them. It does so, even if the switch takes place with messages for the queue Q1 at QM3 still on SYSTEM.CLUSTER.TRANSMIT.QUEUE.

The configuration to isolate cluster message traffic in Figure 131 on page 955 does not result in as great an isolation of traffic as the configuration using remote queues in “Adding a remote queue definition to isolate messages sent from a gateway queue manager” on page 952. If the queue manager QM3 in CL2 is hosting a number of different cluster queues and server applications, all those queues share the cluster

channel, CL2.QM3, connecting QM1 to QM3. The additional flows are illustrated in Figure 131 by the gray arrow representing potential cluster message traffic from the SYSTEM.CLUSTER.TRANSMIT.QUEUE to the cluster-sender channel CL2.QM3.

The remedy is to restrict the queue manager to hosting one cluster queue in a particular cluster. If the queue manager is already hosting a number of cluster queues, then to meet this restriction, you must either create another queue manager, or create another cluster; see “Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 957.

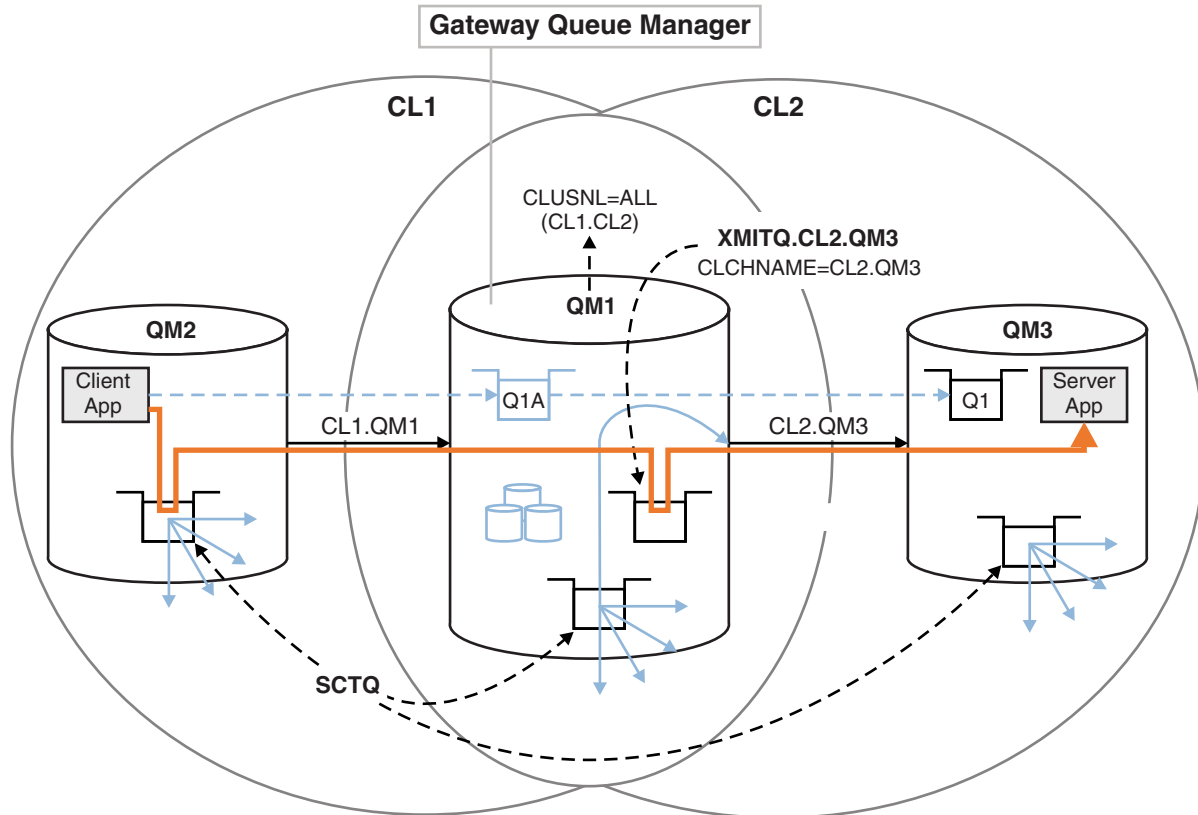


Figure 131. Client-server application deployed to hub and spoke architecture using an additional cluster transmission queue.

### Procedure

1. Create an additional cluster transmission queue for the cluster-sender channel CL2.QM3 on the gateway queue manager, QM1.

```
*... on QM1
DEFINE QLOCAL(XMITQ.CL2.QM3) USAGE(XMITQ) CLCHNAME(CL2.QM3)
```

2. Switch to using the transmission queue, XMITQ.CL2.QM3.

- a. Stop the cluster-sender channel CL2.QM3.

```
*... On QM1
STOP CHANNEL(CL2.QM3)
```

The response is that the command is accepted:

```
AMQ8019: Stop IBM MQ channel accepted.
```

- b. Check that the channel CL2.QM3 is stopped

If the channel does not stop, you can run the **STOP CHANNEL** command again with the **FORCE** option. An example of setting the **FORCE** option would be if the channel does not stop, and you cannot restart the other queue manager to synchronize the channel.

```
*... On QM1
start
```

The response is a summary of the channel status

```
AMQ8417: Display Channel Status details.
CHANNEL (CL2.QM3) CHLTYPE (CLUSSDR)
CONNAME (127.0.0.1(1413)) CURRENT
RQMNAME (QM3) STATUS (STOPPED)
SUBSTATE (MQGET) XMITQ (SYSTEM.CLUSTER.TRANSMIT.QUEUE)
```

- c. Start the channel, CL2.QM3.

```
*... On QM1
START CHANNEL (CL2.QM3)
```

The response is that the command is accepted:

```
AMQ8018: Start IBM MQ channel accepted.
```

- d. Check the channel started.

```
*... On QM1
DISPLAY CHSTATUS (CL2.QM3)
```

The response is a summary of the channel status:

```
AMQ8417: Display Channel Status details.
CHANNEL (CL2.QM3) CHLTYPE (CLUSSDR)
CONNAME (127.0.0.1(1413)) CURRENT
RQMNAME (QM3) STATUS (RUNNING)
SUBSTATE (MQGET) XMITQ (XMITQ.CL2.QM3)
```

- e. Check the transmission queue was switched.

Monitor the gateway queue manager error log for the message “ AMQ7341 The transmission queue for channel CL2.QM3 is XMITQ.CL2.QM3 ”.

## What to do next

Test the separate transmission queue by sending a message from QM2 to Q1 on QM3 using the queue alias definition Q1A

1. Run the sample program **amqsput** on QM2 to put a message.

```
C:\IBM\MQ>amqsput Q1A QM2
Sample AMQSPUT0 start
target queue is Q1A
Sample request message from QM2 to Q1 using Q1A
```

```
Sample AMQSPUT0 end
```

2. Run the sample program **amqsget** to get the message from Q1 on QM3

```
C:\IBM\MQ>amqsget Q1 QM3
Sample AMQSGET0 start
message <Sample request message from QM2 to Q1 using Q1A>
no more messages
Sample AMQSGET0 end
```

**Related concepts:**

“Working with cluster transmission queues and cluster-sender channels” on page 907

Messages between clustered queue managers are stored on cluster transmission queues and forwarded by cluster-sender channels. At any point in time, a cluster-sender channel is associated with one transmission queue. If you change the configuration of the channel, it might switch to a different transmission queue next time it starts. The processing of this switch is automated, and transactional.

**Related tasks:**

“Adding a queue manager to a cluster: separate transmission queues” on page 926

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

**Related information:**

Access control and multiple cluster transmission queues

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

**Adding a cluster and a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager:**

Modify the configuration of overlapping clusters that use a gateway queue manager. After the modification messages are transferred to an application from the gateway queue manager without using the same transmission queue or channels as other cluster messages. The solution uses an additional cluster to isolate the messages to a particular cluster queue.

**Before you begin**

The steps in the task are written to modify the configuration illustrated in Figure 131 on page 955.

1. The gateway queue manager must be on Version 7.5, or later.
2. Construct the overlapping clusters shown in Client-server application deployed to hub and spoke architecture using IBM MQ clusters in “Creating two-overlapping clusters with a gateway queue manager” on page 947 by following the steps in that task.
3. Do the steps in Figure 131 on page 955 in “Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954 to create the solution without the additional cluster. Use this as a base for the steps in this task.

**About this task**

The solution to isolating message traffic to a single application in “Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954 works if the target cluster queue is the only cluster queue on a queue manager. If it is not, you have two choices. Either move the queue to a different queue manager, or create a cluster that isolates the queue from other cluster queues on the queue manager.

This task takes you through the steps to add a cluster to isolate the target queue. The cluster is added just for that purpose. In practice, approach the task of isolating certain applications systematically when you are in the process of designing clusters and cluster naming schemes. Adding a cluster each time a queue requires isolation might end up with many clusters to manage. In this task, you change the configuration in “Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954 by adding a cluster CL3 to isolate Q1 on QM3. Applications continue to run throughout the change.

The new and changed definitions are highlighted in Figure 132 on page 958. The summary of the changes is as follows: Create a cluster, which means you must also create a new full cluster repository. In the example, QM3 is made one of the full repositories for CL3. Create cluster-sender and cluster-receiver channels for QM1 to add the gateway queue manager to the new cluster. Change the definition of Q1 to

switch it to CL3. Modify the cluster namelist on the gateway queue manager, and add a cluster transmission queue to use the new cluster channel. Finally, switch the queue alias Q1A to the new cluster namelist.

IBM MQ cannot transfer messages from the transmission queue XMITQ.CL2.QM3 that you added in “Adding a cluster transmit queue to isolate cluster message traffic sent from a gateway queue manager” on page 954 to the new transmission queue XMITQ.CL3.QM3, automatically. It can transfer messages automatically only if both transmission queues are served by the same cluster-sender channel. Instead, the task describes one way to perform the switch manually, which might be appropriate to you. When the transfer is completed, you have the option of reverting to using the default cluster transmission queue for other CL2 cluster queues on QM3. Or you can continue to use XMITQ.CL2.QM3. If you decide to revert to a default cluster transmission queue, the gateway queue manager manages the switch for you automatically.

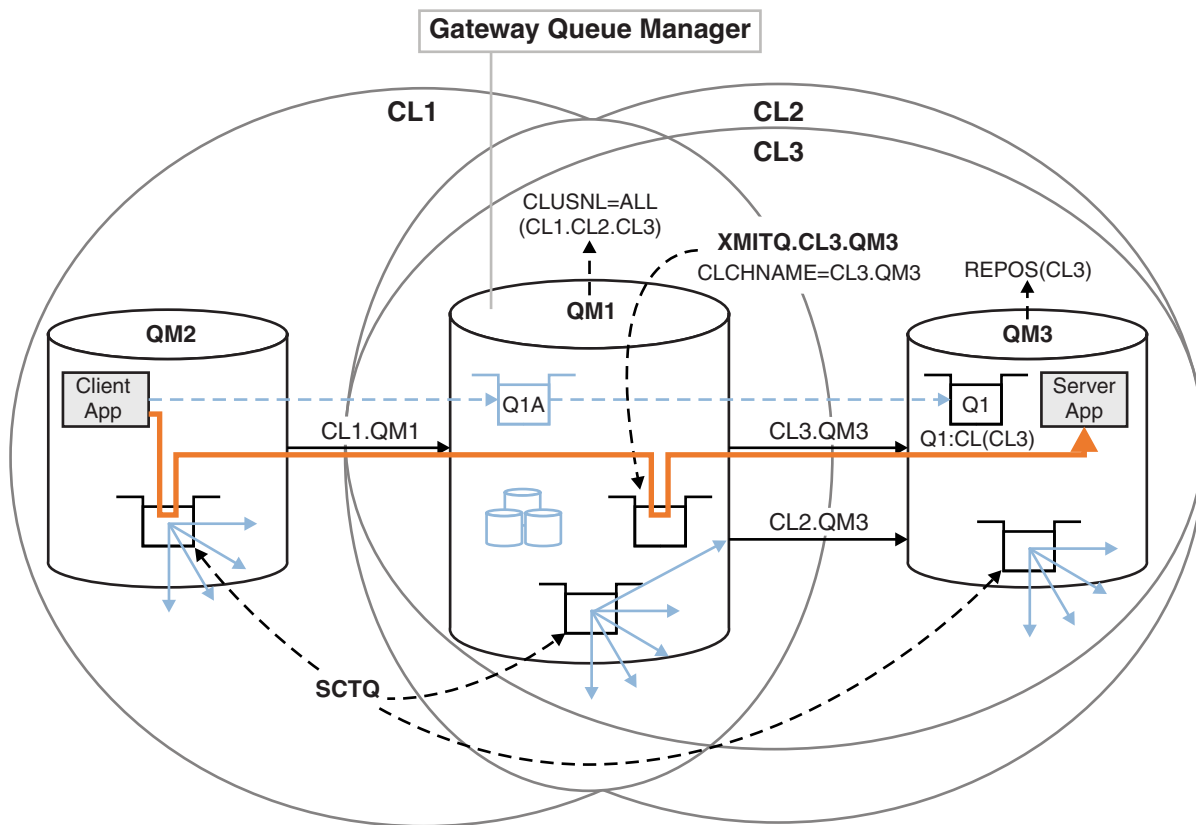


Figure 132. Using an additional cluster to separate message traffic in the gateway queue manager that goes to one of a number of cluster queues on the same queue manager

### Procedure

1. Alter the queue managers QM3 and QM5 to make them repositories for both CL2 and CL3.

To make a queue manager a member of multiple clusters, it must use a cluster name list to identify the clusters it is a member of.

```
*... On QM3 and QM5
DEFINE NAMLIST(CL23) NAMES(CL2, CL3) REPLACE
ALTER QMGR REPOS(' ') REPOSNL(CL23)
```

2. Define the channels between the queue managers QM3 and QM5 for CL3.

```
*... On QM3
DEFINE CHANNEL(CL3.QM5) CHLTYPE(CLUSSDR) CONNAME('localhost(1415)') CLUSTER(CL3) REPLACE
DEFINE CHANNEL(CL3.QM3) CHLTYPE(CLUSRCVR) CONNAME('localhost(1413)') CLUSTER(CL3) REPLACE
```

```
*... On QM5
DEFINE CHANNEL(CL3.QM3) CHLTYPE(CLUSSDR) CONNAME('localhost(1413)') CLUSTER(CL3) REPLACE
DEFINE CHANNEL(CL3.QM5) CHLTYPE(CLUSRCVR) CONNAME('localhost(1415)') CLUSTER(CL3) REPLACE
```

3. Add the gateway queue manager to CL3.

Add the gateway queue manager by adding QM1 to CL3 as a partial repository. Create a partial repository by adding cluster-sender and cluster-receiver channels to QM1.

Also, add CL3 to the name list of all clusters connected to the gateway queue manager.

```
*... On QM1
DEFINE CHANNEL(CL3.QM3) CHLTYPE(CLUSSDR) CONNAME('localhost(1413)') CLUSTER(CL3) REPLACE
DEFINE CHANNEL(CL3.QM1) CHLTYPE(CLUSRCVR) CONNAME('localhost(1411)') CLUSTER(CL3) REPLACE
ALTER NAMLIST(ALL) NAMES(CL1, CL2, CL3)
```

4. Add a cluster transmission queue to the gateway queue manager, QM1, for messages going to CL3 on QM3.

Initially, stop the cluster-sender channel transferring messages from the transmission queue until you are ready to switch transmission queues.

```
*... On QM1
DEFINE QLOCAL(XMITQ.CL3.QM3) USAGE(XMITQ) CLCHNAME(CL3.QM3) GET(DISABLED) REPLACE
```

5. Drain messages from the existing cluster transmission queue XMITQ.CL2.QM3.

This subprocedure is intended to preserve the order of messages in Q1 to match the order they arrived at the gateway queue manager. With clusters, message ordering is not fully guaranteed, but is likely. If guaranteed message ordering is required, applications must define the order of messages; see The order in which messages are retrieved from a queue.

- a. Change the target queue Q1 on QM3 from CL2 to CL3.

```
*... On QM3
ALTER QLOCAL(Q1) CLUSTER(CL3)
```

- b. Monitor XMITQ.CL3.QM3 until messages start to be delivered to it.

Messages start to be delivered to XMITQ.CL3.QM3 when the switch of Q1 to CL3 is propagated to the gateway queue manager.

```
*... On QM1
DISPLAY QUEUE(XMITQ.CL3.QM3) CURDEPTH
```

- c. Monitor XMITQ.CL2.QM3 until it has no messages waiting to be delivered to Q1 on QM3.

**Note:** XMITQ.CL2.QM3 might be storing messages for other queues on QM3 that are members of CL2, in which case the depth might not go to zero.

```
*... On QM1
DISPLAY QUEUE(XMITQ.CL2.QM3) CURDEPTH
```

- d. Enable get from the new cluster transmission queue, XMITQ.CL3.QM3

```
*... On QM1
ALTER QLOCAL(XMITQ.CL3.QM3) GET(ENABLED)
```

6. Remove the old cluster transmission queue, XMITQ.CL2.QM3, if it is no longer required.

Messages for cluster queues in CL2 on QM3 revert to using the default cluster transmission queue on the gateway queue manager, QM1. The default cluster transmission queue is either SYSTEM.CLUSTER.TRANSMIT.QUEUE, or SYSTEM.CLUSTER.TRANSMIT.CL2.QM3. Which one depends on whether the value of the queue manager attribute **DEFCLXQ** on QM1 is SCTQ or CHANNEL. The queue manager transfers messages from XMITQ.CL2.QM3 automatically when the cluster-sender channel CL2.QM3 next starts.

- a. Change the transmission queue, XMITQ.CL2.QM3, from being a cluster transmission queue to being a normal transmission queue.

This breaks the association of the transmission queue with any cluster-sender channels. In response, IBM MQ automatically transfers messages from XMITQ.CL2.QM3 to the default cluster transmission queue when the cluster-sender channel is next started. Until then, messages for CL2 on QM3 continue to be placed on XMITQ.CL2.QM3.

```
*... On QM1
ALTER QLOCAL(XMITQ.CL2.QM3) CLCHNAME(' ')
```

- b. Stop the cluster-sender channel CL2.QM3.

Stopping and restarting the cluster-sender channel initiates the transfer of messages from XMITQ.CL2.QM3 to the default cluster transmission queue. Typically you would stop and start the channel manually to start the transfer. The transfer starts automatically if the channel restarts after shutting down on the expiry of its disconnect interval.

```
*... On QM1
STOP CHANNEL(CL2.QM3)
```

The response is that the command is accepted:

```
AMQ8019: Stop IBM MQ channel accepted.
```

- c. Check that the channel CL2.QM3 is stopped

If the channel does not stop, you can run the **STOP CHANNEL** command again with the **FORCE** option. An example of setting the **FORCE** option would be if the channel does not stop, and you cannot restart the other queue manager to synchronize the channel.

```
*... On QM1
DISPLAY CHSTATUS(CL2.QM3)
```

The response is a summary of the channel status

```
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1413)) CURRENT
RQMNAME(QM3) STATUS(STOPPED)
SUBSTATE(MQGET) XMITQ(XMITQ.CL2.QM3)
```

- d. Start the channel, CL2.QM3.

```
*... On QM1
START CHANNEL(CL2.QM3)
```

The response is that the command is accepted:

```
AMQ8018: Start IBM MQ channel accepted.
```

- e. Check the channel started.

```
*... On QM1
DISPLAY CHSTATUS(CL2.QM3)
```

The response is a summary of the channel status:

```
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1413)) CURRENT
RQMNAME(QM3) STATUS(RUNNING)
SUBSTATE(MQGET) XMITQ(SYSTEM.CLUSTER.TRANSMIT. QUEUE|CL2.QM3)
```

- f. Monitor the gateway queue manager error log for the message " AMQ7341 The transmission queue for channel CL2.QM3 is SYSTEM.CLUSTER.TRANSMIT. QUEUE|CL2.QM3 ".

- g. Delete the cluster transmission queue, XMITQ.CL2.QM3.

```
*... On QM1
DELETE QLOCAL(XMITQ.CL2.QM3)
```

## What to do next

Test the separately clustered queue by sending a message from QM2 to Q1 on QM3 using the queue alias definition Q1A

1. Run the sample program **amqspu**t on QM2 to put a message.



```
C:\IBM\MQ>amqsput Q1A QM2
Sample AMQSPUT0 start
target queue is Q1A
Sample request message from QM2 to Q1 using Q1A
```

```
Sample AMQSPUT0 end
```

2. Run the sample program **amqsget** to get the message from Q1 on QM3

```
C:\IBM\MQ>amqsget Q1 QM3
Sample AMQSGET0 start
message <Sample request message from QM2 to Q1 using Q1A>
no more messages
Sample AMQSGET0 end
```

### Related concepts:

“Working with cluster transmission queues and cluster-sender channels” on page 907

Messages between clustered queue managers are stored on cluster transmission queues and forwarded by cluster-sender channels. At any point in time, a cluster-sender channel is associated with one transmission queue. If you change the configuration of the channel, it might switch to a different transmission queue next time it starts. The processing of this switch is automated, and transactional.

### Related tasks:

“Adding a queue manager to a cluster: separate transmission queues” on page 926

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

### Related information:

Access control and multiple cluster transmission queues

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

### Changing the default to separate cluster transmission queues to isolate message traffic:

You can change the default way a queue manager stores messages for a clustered queue or topic on a transmission queue. Changing the default provides you with a way to isolate cluster messages on a gateway queue manager.

### Before you begin

1. The gateway queue manager must be on Version 7.5, or later.
2. Construct the overlapping clusters shown in Client-server application deployed to hub and spoke architecture using IBM MQ clusters in “Creating two-overlapping clusters with a gateway queue manager” on page 947 by following the steps in that task.

### About this task

To implement the architecture with multiple clusters queue, your gateway queue manager must be on Version 7.5, or later. All you do to use multiple cluster transmission queues is to change the default cluster transmission queue type on the gateway queue manager. Change the value of the queue manager attribute **DEFCLXQ** on QM1 from SCTQ to CHANNEL ; see Figure 133 on page 962. The diagram shows one message flow. For flows to other queue managers, or to other clusters, the queue manager creates additional permanent dynamic cluster transmission queues. Each cluster-sender channel transfers messages from a different cluster transmission queue.

The change does not take effect immediately, unless you are connecting the gateway queue manager to clusters for the first time. The task includes steps for the typical case of managing a change to an existing configuration. To set up a queue manager to use separate cluster transmission queues when it first joins a

cluster; see “Adding a queue manager to a cluster: separate transmission queues” on page 926.

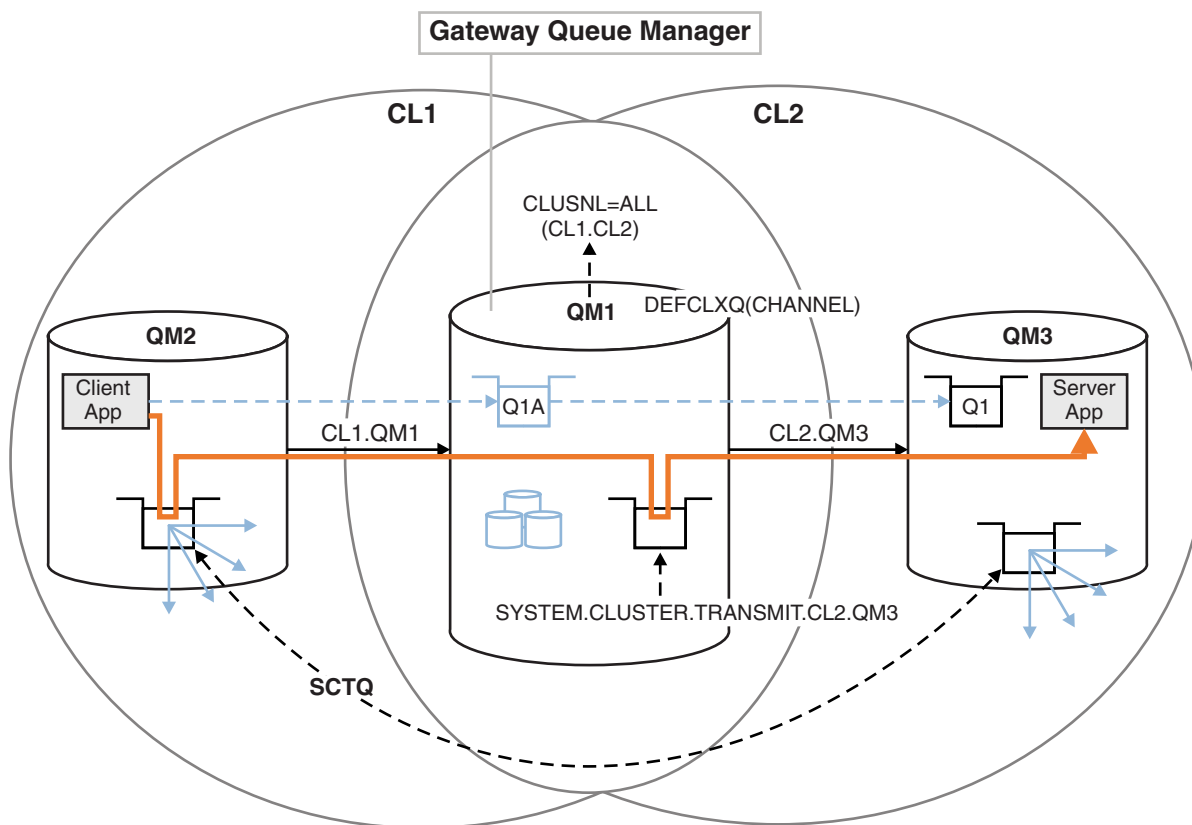


Figure 133. Client-server application deployed to hub and spoke architecture with separate cluster transmission queues on the gateway queue manager.

### Procedure

1. Change the gateway queue manager to use separate cluster transmission queues.

```
*... On QM1
ALTER QMGR DEFCLXQ(CHANNEL)
```

2. Switch to the separate cluster transmission queues.

Any cluster-sender channel that is not running switches to using separate cluster transmission queues when it next starts.

To switch the running channels, either restart the queue manager, or follow these steps:

- a. List the cluster-sender channels that are running with SYSTEM.CLUSTER.TRANSMIT.QUEUE.

```
*... On QM1
DISPLAY CHSTATUS(*) WHERE(XMITQ EQ 'SYSTEM.CLUSTER.TRANSMIT.QUEUE')
```

The response is list of channel status reports:

```
AMQ8417: Display Channel Status details.
CHANNEL(CL1.QM2) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1412)) CURRENT
RQMNAME(QM2) STATUS(RUNNING)
SUBSTATE(MQGET) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1413)) CURRENT
RQMNAME(QM3) STATUS(RUNNING)
```

```

SUBSTATE(MQGET) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM5) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1415)) CURRENT
RQMNAME(QM5) STATUS(RUNNING)
SUBSTATE(MQGET) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL1.QM4) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1414)) CURRENT
RQMNAME(QM4) STATUS(RUNNING)
SUBSTATE(MQGET) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)

```

b. Stop the channels that are running

For each channel in the list, run the command:

```

*... On QM1
STOP CHANNEL(ChannelName)

```

Where *ChannelName* is each of CL1.QM2, CL1.QM4, CL1.QM3, CL1.QM5.

The response is that the command is accepted:

```

AMQ8019: Stop IBM MQ channel accepted.

```

c. Monitor which channels are stopped

```

*... On QM1
DISPLAY CHSTATUS(*) WHERE(XMITQ EQ 'SYSTEM.CLUSTER.TRANSMIT.QUEUE')

```

The response is a list of channels that are still running and channels that are stopped:

```

AMQ8417: Display Channel Status details.
CHANNEL(CL1.QM2) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1412)) CURRENT
RQMNAME(QM2) STATUS(STOPPED)
SUBSTATE() XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1413)) CURRENT
RQMNAME(QM3) STATUS(STOPPED)
SUBSTATE() XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL2.QM5) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1415)) CURRENT
RQMNAME(QM5) STATUS(STOPPED)
SUBSTATE() XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
AMQ8417: Display Channel Status details.
CHANNEL(CL1.QM4) CHLTYPE(CLUSSDR)
CONNNAME(127.0.0.1(1414)) CURRENT
RQMNAME(QM4) STATUS(STOPPED)
SUBSTATE() XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)

```

d. Start each stopped channel.

Do this step for all the channels that were running. If a channel does not stop, you can run the **STOP CHANNEL** command again with the **FORCE** option. An example of setting the **FORCE** option would be if the channel does not stop, and you cannot restart the other queue manager to synchronize the channel.

```

*... On QM1
START CHANNEL(CL2.QM5)

```

The response is that the command is accepted:

AMQ8018: Start IBM MQ channel accepted.

- e. Monitor the transmission queues being switched.

Monitor the gateway queue manager error log for the message " AMQ7341 The transmission queue for channel CL2.QM3 is SYSTEM.CLUSTER.TRANSMIT.QUEUE|CL2.QM3 ".

- f. Check that SYSTEM.CLUSTER.TRANSMIT.QUEUE is no longer used

```
*... On QM1
DISPLAY CHSTATUS(*) WHERE(XMITQ EQ 'SYSTEM.CLUSTER.TRANSMIT.QUEUE')
DISPLAY QUEUE(SYSTEM.CLUSTER.TRANSMIT.QUEUE) CURDEPTH
```

The response is list of channel status reports, and the depth of SYSTEM.CLUSTER.TRANSMIT.QUEUE:

AMQ8420: Channel Status not found.

AMQ8409: Display Queue details.

```
QUEUE(SYSTEM.CLUSTER.TRANSMIT.QUEUE) TYPE(QLOCAL)
CURDEPTH(0)
```

- g. Monitor which channels are started

```
*... On QM1
DISPLAY CHSTATUS(*) WHERE(XMITQ LK 'SYSTEM.CLUSTER.TRANSMIT.*')
```

The response is a list of the channels, in this case already running with the new default cluster transmission queues:

AMQ8417: Display Channel Status details.

```
CHANNEL(CL1.QM2) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1412)) CURRENT
RQMNAME(QM2) STATUS(RUNNING)
SUBSTATE(MQGET)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.CL1.QM2)
```

AMQ8417: Display Channel Status details.

```
CHANNEL(CL2.QM3) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1413)) CURRENT
RQMNAME(QM3) STATUS(RUNNING)
SUBSTATE(MQGET)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.CL2.QM3)
```

AMQ8417: Display Channel Status details.

```
CHANNEL(CL2.QM5) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1415)) CURRENT
RQMNAME(QM5) STATUS(RUNNING)
SUBSTATE(MQGET)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.CL2.QM5)
```

AMQ8417: Display Channel Status details.

```
CHANNEL(CL1.QM4) CHLTYPE(CLUSSDR)
CONNAME(127.0.0.1(1414)) CURRENT
RQMNAME(QM4) STATUS(RUNNING)
SUBSTATE(MQGET)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.CL1.QM4)
```

## What to do next

1. Test the automatically defined cluster transmission queue by sending a message from QM2 to Q1 on QM3, resolving queue name with the queue alias definition Q1A
  - a. Run the sample program **amqsput** on QM2 to put a message.

```
C:\IBM\MQ>amqsput Q1A QM2
Sample AMQSPUT0 start
target queue is Q1A
```

Sample request message from QM2 to Q1 using Q1A

Sample AMQSPUT0 end

- b. Run the sample program **amqsget** to get the message from Q1 on QM3

```
C:\IBM\MQ>amqsget Q1 QM3
```

```
Sample AMQSGET0 start
```

```
message <Sample request message from QM2 to Q1 using Q1A>
```

```
no more messages
```

```
Sample AMQSGET0 end
```

2. Consider whether to reconfigure security, by configuring security for the cluster queues on the queue managers where messages for the cluster queues originate.

#### **Related tasks:**

“Adding a queue manager to a cluster: separate transmission queues” on page 926

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using multiple cluster transmission queues.

#### **Related information:**

Access control and multiple cluster transmission queues

Clustering: Application isolation using multiple cluster transmission queues

Clustering: Planning how to configure cluster transmission queues

### **Removing a cluster queue from a queue manager**

Disable the INVENTQ queue at Toronto. Send all the inventory messages to New York, and delete the INVENTQ queue at Toronto when it is empty.

#### **Before you begin**

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager that hosts a queue” on page 931. It contains four queue managers. LONDON and NEWYORK both hold full repositories. PARIS and TORONTO hold partial repositories. The inventory application runs on the systems in New York and Toronto and is driven by the arrival of messages on the INVENTQ queue.
- Because of reduced workload, you no longer want to run the inventory application in Toronto. You want to disable the INVENTQ queue hosted by the queue manager TORONTO, and have TORONTO feed messages to the INVENTQ queue in NEWYORK.
- Network connectivity exists between all four systems.
- The network protocol is TCP.

#### **About this task**

Follow these steps to remove a cluster queue.

#### **Procedure**

1. Indicate that the queue is no longer available.

To remove a queue from a cluster, remove the cluster name from the local queue definition. Alter the INVENTQ on TORONTO so that it is not accessible from the rest of the cluster:

```
ALTER QLOCAL(INVENTQ) CLUSTER(' ')
```

2. Check that the queue is no longer available.

On a full repository queue manager, either LONDON or NEWYORK, check that the queue is no longer hosted by queue manager TORONTO by issuing the following command:

```
DIS QCLUSTER (INVENTQ)
```

TORONTO is not listed in the results, if the ALTER command has completed successfully.

3. Disable the queue.

Disable the INVENTQ queue at TORONTO so that no further messages can be written to it:

```
ALTER QLOCAL(INVENTQ) PUT(DISABLED)
```

Now messages in transit to this queue using MQ00\_BIND\_ON\_OPEN go to the dead-letter queue. You need to stop all applications from putting messages explicitly to the queue on this queue manager.

4. Monitor the queue until it is empty.

Monitor the queue using the DISPLAY QUEUE command, specifying the attributes IPPROCS, OPPROCS, and CURDEPTH, or use the **WRKMQMSTS** command on IBM i. When the number of input and output processes, and the current depth of the queues are all zero, the queue is empty.

5. Monitor the channel to ensure there are no in-doubt messages.

To be sure that there are no in-doubt messages on the channel INVENTORY.TORONTO, monitor the cluster-sender channel called INVENTORY.TORONTO on each of the other queue managers. Issue the DISPLAY CHSTATUS command specifying the INDOUBT parameter from each queue manager:

```
DISPLAY CHSTATUS(INVENTORY.TORONTO) INDOUBT
```

If there are any in-doubt messages, you must resolve them before proceeding. For example, you might try issuing the RESOLVE channel command or stopping and restarting the channel.

6. Delete the local queue.

When you are satisfied that there are no more messages to be delivered to the inventory application at TORONTO, you can delete the queue:

```
DELETE QLOCAL(INVENTQ)
```

7. You can now remove the inventory application from the system in Toronto Removing the application avoids duplication and saves space on the system.

## Results

The cluster set up by this task is like that setup by the previous task. The difference is the INVENTQ queue is no longer available at queue manager TORONTO.

When you took the queue out of service in step 1, the TORONTO queue manager sent a message to the two full repository queue managers. It notified them of the change in status. The full repository queue managers pass on this information to other queue managers in the cluster that have requested updates to information concerning the INVENTQ.

When a queue manager puts a message on the INVENTQ queue the updated partial repository indicates that the INVENTQ queue is available only at the NEWYORK queue manager. The message is sent to the NEWYORK queue manager.

## What to do next

In this task, there was only one queue to remove and only one cluster to remove it from.

Suppose that there are many queues referring to a namelist containing many cluster names. For example, the TORONTO queue manager might host not only the INVENTQ, but also the PAYROLLQ, SALESQ, and PURCHASESQ. TORONTO makes these queues available in all the appropriate clusters, INVENTORY, PAYROLL, SALES, and PURCHASES. Define a namelist of the cluster names on the TORONTO queue manager:

```
DEFINE NAMELIST(TOROLIST)
DESCR('List of clusters TORONTO is in')
NAMES(INVENTORY, PAYROLL, SALES, PURCHASES)
```

Add the namelist to each queue definition:

```
DEFINE QLOCAL(INVENTQ) CLUSNL(TOROLIST)
DEFINE QLOCAL(PAYROLLQ) CLUSNL(TOROLIST)
DEFINE QLOCAL(SALESQ) CLUSNL(TOROLIST)
DEFINE QLOCAL(PURCHASESQ) CLUSNL(TOROLIST)
```

Now suppose that you want to remove all those queues from the SALES cluster, because the SALES operation is to be taken over by the PURCHASES operation. All you need to do is alter the TOROLIST namelist to remove the name of the SALES cluster from it.

If you want to remove a single queue from one of the clusters in the namelist, create a namelist, containing the remaining list of cluster names. Then alter the queue definition to use the new namelist. To remove the PAYROLLQ from the INVENTORY cluster:

1. Create a namelist:

```
DEFINE NAMELIST(TOROSHORTLIST)
DESCR('List of clusters TORONTO is in other than INVENTORY')
NAMES(PAYROLL, SALES, PURCHASES)
```

2. Alter the PAYROLLQ queue definition:

```
ALTER QLOCAL(PAYROLLQ) CLUSNL(TOROSHORTLIST)
```

## Removing a queue manager from a cluster

Remove a queue manager from a cluster, in scenarios where the queue manager can communicate normally with at least one full repository in the cluster.

### Before you begin

This method is the best practice for scenarios in which at least one full repository is available, and can be contacted by the queue manager that is being removed. This method involves the least manual intervention, and allows the queue manager to negotiate a controlled withdrawal from the cluster. If the queue manager that is being removed cannot contact a full repository, see “Removing a queue manager from a cluster: Alternative method” on page 969.

Before you remove the queue manager from the cluster, you must ensure that the queue manager is no longer hosting resources that are needed by the cluster:

- If the queue manager hosts a full repository, complete steps 1-6 from “Moving a full repository to another queue manager” on page 935. If the full repository functionality of the queue manager to be removed is not to be moved to a different queue manager, it is only necessary to complete steps 5 and 6.
- If the queue manager hosts cluster queues, complete steps 1-7 from “Removing a cluster queue from a queue manager” on page 965.
- If the queue manager hosts cluster topics, either delete the topics (for example by using the DELETE TOPIC command), or move them to other hosts as described in “Moving a cluster topic definition to a different queue manager” on page 1011.

**Note:** If you remove a queue manager from a cluster, and the queue manager still hosts a cluster topic, then the queue manager might continue to attempt to deliver publications to the queue managers that are left in the cluster until the topic is deleted.

### About this task

This example task removes the queue manager LONDON from the INVENTORY cluster. The INVENTORY cluster is set up as described in “Adding a queue manager to a cluster” on page 924, and modified as described in “Removing a cluster queue from a queue manager” on page 965.

The process for removing a queue manager from a cluster is more complicated than the process of adding a queue manager.

When a queue manager joins a cluster, the existing members of the cluster have no knowledge of the new queue manager and so have no interactions with it. New sender and receiver channels must be created on the joining queue manager so that it can connect to a full repository.

When a queue manager is removed from a cluster, it is likely that applications connected to the queue manager are using objects such as queues that are hosted elsewhere in the cluster. Also, applications that are connected to other queue managers in the cluster might be using objects hosted on the target queue manager. As a result of these applications, the current queue manager might create additional sender channels to establish communication with cluster members other than the full repository that it used to join the cluster. Every queue manager in the cluster has a cached copy of data that describes other cluster members. This might include the one that is being removed.

## Procedure

1. Alter the manually defined cluster receiver channels to remove them from the cluster, on queue manager LONDON:

```
ALTER CHANNEL(INVENTORY.LONDON) CHLTYPE(CLUSRCVR) CLUSTER(' ')
```

2. Alter the manually defined cluster sender channels to remove them from the cluster, on queue manager LONDON:

```
ALTER CHANNEL(INVENTORY.PARIS) CHLTYPE(CLUSSDR) CLUSTER(' ')
```

The other queue managers in the cluster learn that this queue manager and its cluster resources are no longer part of the cluster.

3. Monitor the cluster transmit queue, on queue manager LONDON, until there are no messages that are waiting to flow to any full repository in the cluster.

```
DISPLAY CHSTATUS(INVENTORY.LONDON) XQMSGSA
```

If messages remain on the transmit queue, determine why they are not being sent to the PARIS and NEWYORK full repositories before continuing.

## Results

The queue manager LONDON is no longer part of the cluster. However, it can still function as an independent queue manager.

## What to do next

The result of these changes can be confirmed by issuing the following command on the remaining members of the cluster:

```
DISPLAY CLUSQMGR(LONDON)
```

The queue manager continues to be displayed until the auto-defined cluster sender channels to it have stopped. You can wait for this to happen, or, continue to monitor for active instances by issuing the following command:

```
DISPLAY CHANNEL(INVENTORY.LONDON)
```

When you are confident that no more messages are being delivered to this queue manager, you can stop the cluster sender channels to LONDON by issuing the following command on the remaining members of the cluster:

```
STOP CHANNEL(INVENTORY.LONDON) STATUS(INACTIVE)
```

After the changes are propagated throughout the cluster, and no more messages are being delivered to this queue manager, stop and delete the CLUSRCVR channel on LONDON:



```
STOP CHANNEL(INVENTORY.LONDON)
DELETE CHANNEL(INVENTORY.LONDON)
```

The removed queue manager can be added back into the cluster at a later point as described in “Adding a queue manager to a cluster” on page 924. The removed queue manager continues to cache knowledge of the remaining members of the cluster for up to 90 days. If you prefer not to wait until this cache expires, it can be forcibly removed as described in “Restoring a queue manager to its pre-cluster state” on page 971.

### **Removing a queue manager from a cluster: Alternative method:**

Remove a queue manager from a cluster, in scenarios where, because of a significant system or configuration issue, the queue manager cannot communicate with any full repository in the cluster.

#### **Before you begin**

This alternative method of removing a queue manager from a cluster manually stops and deletes all cluster channels linking the removed queue manager to the cluster, and forcibly removes the queue manager from the cluster. This method is used in scenarios where the queue manager that is being removed cannot communicate with any of the full repositories. This might be (for example) because the queue manager has stopped working, or because there has been a prolonged communications failure between the queue manager and the cluster. Otherwise, use the most common method: “Removing a queue manager from a cluster” on page 967.

Before you remove the queue manager from the cluster, you must ensure that the queue manager is no longer hosting resources that are needed by the cluster:

- If the queue manager hosts a full repository, complete steps 1-6 from “Moving a full repository to another queue manager” on page 935. If the full repository functionality of the queue manager to be removed is not to be moved to a different queue manager, it is only necessary to complete steps 5 and 6.
- If the queue manager hosts cluster queues, complete steps 1-7 from “Removing a cluster queue from a queue manager” on page 965.
- If the queue manager hosts cluster topics, either delete the topics (for example by using the DELETE TOPIC command), or move them to other hosts as described in “Moving a cluster topic definition to a different queue manager” on page 1011.

**Note:** If you remove a queue manager from a cluster, and the queue manager still hosts a cluster topic, then the queue manager might continue to attempt to deliver publications to the queue managers that are left in the cluster until the topic is deleted.

#### **About this task**

This example task removes the queue manager LONDON from the INVENTORY cluster. The INVENTORY cluster is set up as described in “Adding a queue manager to a cluster” on page 924, and modified as described in “Removing a cluster queue from a queue manager” on page 965.

The process for removing a queue manager from a cluster is more complicated than the process of adding a queue manager.

When a queue manager joins a cluster, the existing members of the cluster have no knowledge of the new queue manager and so have no interactions with it. New sender and receiver channels must be created on the joining queue manager so that it can connect to a full repository.

When a queue manager is removed from a cluster, it is likely that applications connected to the queue manager are using objects such as queues that are hosted elsewhere in the cluster. Also, applications that are connected to other queue managers in the cluster might be using objects hosted on the target queue

manager. As a result of these applications, the current queue manager might create additional sender channels to establish communication with cluster members other than the full repository that it used to join the cluster. Every queue manager in the cluster has a cached copy of data that describes other cluster members. This might include the one that is being removed.

This procedure might be appropriate in an emergency, when it is not possible to wait for the queue manager to leave the cluster gracefully.

### Procedure

1. Stop all channels used to communicate with other queue managers in the cluster. Use `MODE(FORCE)` to stop the `CLUSRCVR` channel, on queue manager `LONDON`. Otherwise you might need to wait for the sender queue manager to stop the channel:

```
STOP CHANNEL(INVENTORY.LONDON) MODE(FORCE)
STOP CHANNEL(INVENTORY.TORONTO)
STOP CHANNEL(INVENTORY.PARIS)
STOP CHANNEL(INVENTORY.NEWYORK)
```

2. Monitor the channel states, on queue manager `LONDON`, until the channels stop:

```
DISPLAY CHSTATUS(INVENTORY.LONDON)
DISPLAY CHSTATUS(INVENTORY.TORONTO)
DISPLAY CHSTATUS(INVENTORY.PARIS)
DISPLAY CHSTATUS(INVENTORY.NEWYORK)
```

No more application messages are sent to or from the other queue managers in the cluster after the channels stop.

3. Delete the manually defined cluster channels, on queue manager `LONDON`:

```
DELETE CHANNEL(INVENTORY.NEWYORK)
DELETE CHANNEL(INVENTORY.TORONTO)
```

4. The remaining queue managers in the cluster still retain knowledge of the removed queue manager, and might continue to send messages to it. To purge the knowledge from the remaining queue managers, reset the removed queue manager from the cluster on one of the full repositories:

```
RESET CLUSTER(INVENTORY) ACTION(FORCEREMOVE) QMNAME(LONDON) QUEUES(YES)
```

If there might be another queue manager in the cluster that has the same name as the removed queue manager, specify the **QMID** of the removed queue manager.

### Results

The queue manager `LONDON` is no longer part of the cluster. However, it can still function as an independent queue manager.

### What to do next

The result of these changes can be confirmed by issuing the following command on the remaining members of the cluster:

```
DISPLAY CLUSQMGR(LONDON)
```

The queue manager continues to be displayed until the auto-defined cluster sender channels to it have stopped. You can wait for this to happen, or, continue to monitor for active instances by issuing the following command:

```
DISPLAY CHANNEL(INVENTORY.LONDON)
```

After the changes are propagated throughout the cluster, and no more messages are being delivered to this queue manager, delete the `CLUSRCVR` channel on `LONDON`:

```
DELETE CHANNEL(INVENTORY.LONDON)
```

The removed queue manager can be added back into the cluster at a later point as described in “Adding a queue manager to a cluster” on page 924. The removed queue manager continues to cache knowledge of the remaining members of the cluster for up to 90 days. If you prefer not to wait until this cache expires, it can be forcibly removed as described in “Restoring a queue manager to its pre-cluster state.”

## Restoring a queue manager to its pre-cluster state

When a queue manager is removed from a cluster, it retains knowledge of the remaining cluster members. This knowledge eventually expires and is deleted automatically. However, if you prefer to delete it immediately, you can use the steps in this topic.

### Before you begin

It is assumed that the queue manager has been removed from the cluster, and is no longer performing any work in the cluster. For example, its queues are no longer receiving messages from the cluster, and no applications are waiting for messages to arrive in these queues.

### About this task

When a queue manager is removed from a cluster, it retains knowledge of the remaining cluster members for up to 90 days. This can have system benefits, particularly if the queue manager quickly rejoins the cluster. When this knowledge eventually expires, it is deleted automatically. However, there are reasons why you might prefer to delete this information manually. For example:

- You might want to confirm that you have stopped every application on this queue manager that previously used cluster resources. Until the knowledge of the remaining cluster members expires, any such application continues to write to a transmit queue. After the cluster knowledge is deleted, the system generates an error message when such an application tries to use cluster resources.
- When you display status information for the queue manager, you might prefer not to see expiring information about remaining cluster members.

This task uses the INVENTORY cluster as an example. The LONDON queue manager has been removed from the INVENTORY cluster as described in “Removing a queue manager from a cluster” on page 967. To delete knowledge of the remaining members of the cluster, issue the following commands on the LONDON queue manager.

### Procedure

1. Remove all memory of the other queue managers in the cluster from this queue manager:

```
REFRESH CLUSTER(INVENTORY) REPOS(YES)
```

2. Monitor the queue manager until all the cluster resources are gone:

```
DISPLAY CLUSQMGR(*) CLUSTER(INVENTORY)
DISPLAY QCLUSTER(*) CLUSTER(INVENTORY)
DISPLAY TOPIC(*) CLUSTER(INVENTORY)
```

## Related information:

Clusters

Comparison of clustering and distributed queuing

Cluster components

## Maintaining a queue manager

Suspend and resume a queue manager from a cluster to perform maintenance.

### About this task

From time to time, you might need to perform maintenance on a queue manager that is part of a cluster. For example, you might need to take backups of the data in its queues, or apply fixes to the software. If the queue manager hosts any queues, its activities must be suspended. When the maintenance is complete, its activities can be resumed.

### Procedure

1. Suspend a queue manager, by issuing the SUSPEND QMGR **runmqsc** command:

```
SUSPEND QMGR CLUSTER(SALES)
```

The SUSPEND **runmqsc** command notifies the queue managers in the SALES cluster that this queue manager has been suspended.

The purpose of the SUSPEND QMGR command is only to advise other queue managers to avoid sending messages to this queue manager if possible. It does not mean that the queue manager is disabled. Some messages that have to be handled by this queue manager are still sent to it, for example when this queue manager is the only host of a clustered queue.

While the queue manager is suspended the workload management routines avoid sending messages to it. Messages that have to be handled by that queue manager include messages sent by the local queue manager.

IBM MQ uses a workload balancing algorithm to determine which destinations are suitable, rather than selecting the local queue manager whenever possible.

- a. Enforce the suspension of a queue manager by using the FORCE option on the SUSPEND QMGR command:

```
SUSPEND QMGR CLUSTER(SALES) MODE(FORCE)
```

MODE(FORCE) forcibly stops all inbound channels from other queue managers in the cluster. If you do not specify MODE(FORCE), the default MODE(QUIESCE) applies.


2. Do whatever maintenance tasks are necessary.
3. Resume the queue manager by issuing the RESUME QMGR **runmqsc** command:

```
RESUME QMGR CLUSTER(SALES)
```


### Results


The RESUME **runmqsc** command notifies the full repositories that the queue manager is available again. The full repository queue managers disseminate this information to other queue managers that have requested updates to information concerning this queue manager.

## Maintaining the cluster transmission queue


Make every effort to keep cluster transmission queues available. They are essential to the performance of clusters.  On z/OS, set the INDXTYPE of a cluster transmission queue to CORRELID.

### Before you begin

- Make sure that the cluster transmission queue does not become full.
- Take care not to issue an ALTER **runmqsc** command to set it either get disabled or put disabled accidentally.
- Make sure that the medium the cluster transmission queue is stored on  (for example z/OS page sets) does not become full.

 z/OS

### About this task

The following procedure is only applicable to z/OS. 

### Procedure

Set the INDXTYPE of the cluster transmission queue to CORRELID

## Refreshing a cluster queue manager

You can remove auto-defined channels and auto-defined cluster objects from the local repository using the REFRESH CLUSTER command. No messages are lost.

### Before you begin

You might be asked to use the command by your IBM Support Center. Do not use the command without careful consideration. For example, for large clusters use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Clustering: Using REFRESH CLUSTER best practices.

### About this task

A queue manager can make a fresh start in a cluster. In normal circumstances, you do not need to use the REFRESH CLUSTER command.

### Procedure

Issue the REFRESH CLUSTER **MQSC** command from a queue manager to remove auto-defined cluster queue-manager and queue objects from the local repository.

The command only removes objects that refer to other queue managers, it does not remove objects relating to the local queue manager. The command also removes auto-defined channels. It removes channels that do not have messages on the cluster transmission queue and are not attached to a full repository queue manager.

### Results

Effectively, the REFRESH CLUSTER command allows a queue manager to be cold-started with respect to its full repository content. IBM MQ ensures that no data is lost from your queues.

**Related information:**

Clustering: Using REFRESH CLUSTER best practices

**Recovering a cluster queue manager**

Bring the cluster information about a queue manager up to date using the REFRESH CLUSTER `runmqsc` command. Follow this procedure after recovering a queue manager from a point-in-time backup.

**Before you begin**

You have restored a cluster queue manager from a point-in-time backup of a linear log.

**About this task**

To recover a queue manager in a cluster, restore the queue manager, and then bring the cluster information up to date using the REFRESH CLUSTER `runmqsc` command.

**Note:** For large clusters, using the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See Refreshing in a large cluster can affect performance and availability of the cluster.

**Procedure**

Issue the REFRESH CLUSTER command on the restored queue manager for all clusters in which the queue manager participates.

**What to do next**

There is no need to issue the REFRESH CLUSTER command on any other queue manager.

**Related information:**

Clustering: Using REFRESH CLUSTER best practices

**Configuring cluster channels for availability**

Follow good configuration practices to keep cluster channels running smoothly if there are intermittent network stoppages.

**Before you begin**

Clusters relieve you of the need to define channels, but you still need to maintain them. The same channel technology is used for communication between queue managers in a cluster as is used in distributed queuing. To understand about cluster channels, you need to be familiar with matters such as:

- How channels operate
- How to find their status
- How to use channel exits

**About this task**

You might want to give some special consideration to the following points:

**Procedure**

Consider the following points when configuring cluster channels

- Choose values for HBINT or KAINTE on cluster-sender channels and cluster-receiver channels that do not burden the network with lots of heartbeat or keep alive flows. An interval less than about 10 seconds gives false failures, if your network sometimes slows down and introduces delays of this length.
- Set the BATCHHB value to reduce the window for causing a marooned message because it is indoubt on a failed channel. An indoubt batch on a failed channel is more likely to occur if the batch is given longer to fill. If the message traffic along the channel is sporadic with long periods of time between bursts of messages a failed batch is more likely.
- A problem arises if the cluster-sender end of a channel fails and then tries to restart before the heartbeat or keep alive has detected the failure. The channel-sender restart is rejected if the cluster-receiver end of the channel has remained active. To avoid the failure, arrange for the cluster-receiver channel to be terminated and restarted when a cluster-sender channel attempts to restart.

**z/OS On IBM MQ for z/OS**

Control the problem of the cluster-receiver end of the channel has remaining active using the ADOPTMCA and ADOPTCHK parameters on ALTER QMGR.

**On platforms other than z/OS**

Control the problem of the cluster-receiver end of the channel has remaining active using the AdoptNewMCA, AdoptNewMCATimeout, and AdoptNewMCACheck attributes in the qm.ini file or the Windows NT Registry.

## Routing messages to and from clusters

Use queue aliases, queue manager aliases, and remote queue definitions to connect clusters to external queue managers and other clusters.

For details on routing messages to and from clusters, see the following subtopics:

**Related concepts:**

“Queue-manager aliases and clusters” on page 984

Use queue-manager aliases to hide the name of queue managers when sending messages into or out of a cluster, and to workload balance messages sent to a cluster.

“Queue aliases and clusters” on page 987

Use queue aliases to hide the name of a cluster queue, to cluster a queue, adopt different attributes, or adopt different access controls.

“Reply-to queue aliases and clusters” on page 987

A reply-to queue alias definition is used to specify alternative names for reply information. Reply-to queue alias definitions can be used with clusters just the same as in a distributed queuing environment.

**Related tasks:**

“Configuring a queue manager cluster” on page 902

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

“Setting up a new cluster” on page 913

Follow these instructions to set up the example cluster. Separate instructions describe setting up the cluster on TCP/IP, LU 6.2, and with a single transmission queue or multiple transmission queues. Test the cluster works by sending a message from one queue manager to the other.

**Related information:**

Clusters

Comparison of clustering and distributed queuing

Components of a cluster

## Configuring request/reply to a cluster:

Configure a request/reply message path from a queue manager outside a cluster. Hide the inner details of the cluster by using a gateway queue manager as the communication path to and from the cluster.

### Before you begin

Figure 134 shows a queue manager called QM3 that is outside the cluster called DEMO. QM3 could be a queue manager on an IBM MQ product that does not support clusters. QM3 hosts a queue called Q3, which is defined as follows:

```
DEFINE QLOCAL(Q3)
```

Inside the cluster are two queue managers called QM1 and QM2. QM2 hosts a cluster queue called Q2, which is defined as follows:

```
DEFINE QLOCAL(Q2) CLUSTER(DEMO)
```




Figure 134. Putting from a queue manager outside the cluster

### About this task

Follow the advice in the procedure to set up the path for the request and reply messages.



## Procedure

1. Send the request message to the cluster.

Consider how the queue manager that is outside the cluster puts a message to the queue Q2 at QM2, that is inside the cluster. A queue manager outside the cluster must have a QREMOTE definition for each queue in the cluster that it puts messages to.

- a. Define a remote queue for Q2 on QM3.

```
DEFINE QREMOTE(Q2) RNAME(Q2) RQMNAME(QM2) XMITQ(QM1)
```

Because QM3 is not part of a cluster, it must communicate using distributed queuing techniques. Therefore, it must also have a sender-channel and a transmission queue to QM1. QM1 needs a corresponding receiver channel. The channels and transmission queues are not shown explicitly in Figure 134 on page 976.

In the example, an application at QM3 issues an MQPUT call to put a message to Q2. The QREMOTE definition causes the message to be routed to Q2 at QM2 using the sender-channel that is getting messages from the QM1 transmission queue.

2. Receive the reply message from the cluster.

Use a queue manager alias to create a return path for replies to a queue manager outside the cluster. The gateway, QM1, advertises a queue-manager alias for the queue manager that is outside the cluster, QM3. It advertises QM3 to the queue managers inside the cluster by adding the cluster attribute to a queue manager alias definition for QM3. A queue manager alias definition is like a remote queue definition, but with a blank RNAME.

- a. Define a queue manager alias for QM3 on QM1.

```
DEFINE QREMOTE(QM3) RNAME(' ') RQMNAME(QM3) CLUSTER(DEMO)
```

We must consider the choice of name for the transmission queue used to forward replies back from QM1 to QM3. Implicit in the QREMOTE definition, by the omission of the XMITQ attribute, is the name of the transmission queue is QM3. But QM3 is the same name as we expect to advertise to the rest of the cluster using the queue manager alias. IBM MQ does not allow you to give both the transmission queue and the queue manager alias the same name. One solution is to create a transmission queue to forward messages to QM3 with a different name to the queue manager alias.

- b. Provide the transmission queue name in the QREMOTE definition.

```
DEFINE QREMOTE(QM3) RNAME(' ') RQMNAME(QM3) CLUSTER(DEMO) XMITQ(QM3.XMIT)
```

The new queue manager alias couples the new transmission queue called QM3.XMIT with the QM3 queue manager alias. It is a simple and correct solution, but not wholly satisfactory. It has broken the naming convention for transmission queues that they are given the same name as the target queue manager. Are there any alternative solutions that preserve the transmission queue naming convention?

The problem arises because the requester defaulted to passing QM3 as the reply-to queue manager name in the request message that is sent from QM3. The server on QM2 uses the QM3 reply-to queue manager name to address QM3 in its replies. The solution required QM1 to advertise QM3 as the queue manager alias to return reply messages to and prevented QM1 from using QM3 as the name of the transmission queue.

Instead of defaulting to providing QM3 as the reply-to queue manager name, applications on QM3 need to pass a reply-to queue manager alias to QM1 for reply messages. The gateway queue manager QM1 advertises the queue manager alias for replies to QM3 rather than QM3 itself, avoiding the conflict with the name of the transmission queue.

- c. Define a queue manager alias for QM3 on QM1.

```
DEFINE QREMOTE(QM3.ALIAS) RNAME(' ') RQMNAME(QM3) CLUSTER(DEMO)
```

Two changes to the configuration commands are required.

- 1) The QREMOTE at QM1 now advertises our queue manager alias QM3.ALIAS to the rest of the cluster, coupling it to the name of the real queue manager QM3. QM3 is again the name of the transmission queue to send reply queues back to QM3

2) The client application must provide QM3.ALIAS as the name of the reply-to queue manager when it constructs the request message. You can provide QM3.ALIAS to the client application in one of two ways.

- Code QM3.ALIAS in the reply-to queue manager name field constructed by MQPUT in the MQMD. You must do it this way if you are using a dynamic queue for replies.
- Use a reply-to queue alias, Q3.ALIAS, rather than a reply-to queue when providing the reply-to queue name.

```
DEFINE QREMOTE(Q3.ALIAS) RNAME(Q3) RQMNAME(QM3.ALIAS)
```

## What to do next

**Note:** You cannot demonstrate the use of reply-to queue aliases with **AMQSREQ0**. It opens the reply-to queue using the queue name provided in parameter 3, or the default SYSTEM.SAMPLE.REPLY model queue. You need to modify the sample providing another parameter containing the reply-to queue alias to name the reply-to queue manager alias for MQPUT.

### Related tasks:

“Hiding the name of a cluster target queue manager”

Route a message to a cluster queue that is defined on any queue manager in a cluster without naming the queue manager.

*Hiding the name of a cluster target queue manager:*

Route a message to a cluster queue that is defined on any queue manager in a cluster without naming the queue manager.

### Before you begin

- Avoid revealing the names of queue managers that are inside the cluster to queue managers that are outside the cluster.
  - Resolving references to the queue manager hosting a queue inside the cluster removes the flexibility to do workload balancing.
  - It also makes it difficult for you to change a queue manager hosting a queue in the cluster.
  - The alternative is to replace RQMNAME with a queue manager alias provided by the cluster administrator.
  - “Hiding the name of a cluster target queue manager” describes using a queue manager alias to decouple a queue manager outside a cluster from the management of queue managers inside a cluster.
- However, the suggested way to name transmission queues is to give them the name of the target queue manager. The name of the transmission queue reveals the name of a queue manager in the cluster. You have to choose which rule to follow. You might choose to name the transmission queue using either the queue manager name or the cluster name:

#### **Name the transmission queue using the gateway queue manager name**

Disclosure of the gateway queue manager name to queue managers outside a cluster is a reasonable exception to the rule of hiding cluster queue manager names.

#### **Name the transmission queue using the name of the cluster**

If you are not following the convention of naming transmission queues with the name of the target queue manager, use the cluster name.

### About this task

Modify the task “Configuring request/reply to a cluster” on page 976, to hide the name of the target queue manager inside the cluster.

## Procedure

In the example, see Figure 135, define a queue manager alias on the gateway queue manager QM1 called DEMO:

```
DEFINE QREMOTE(DEMO) RNAME(' ') RQMNAME(' ')
```

The QREMOTE definition on QM1 makes the queue manager alias DEMO known to the gateway queue




Figure 135. Putting from a queue manager outside the cluster

manager. QM3, the queue manager outside the cluster, can use the queue manager alias DEMO to send messages to cluster queues on DEMO, rather than having to use an actual queue manager name. If you adopt the convention of using the cluster name to name the transmission queue connecting to a cluster, then the remote queue definition for Q2 becomes:

```
DEFINE QREMOTE(Q2) RNAME(Q2) RQMNAME(DEMO) XMIT(DEMO)
```

## Results

Messages destined for Q2 on DEMO are placed on the DEMO transmission queue. From the transmission queue they are transferred by the sender-channel to the gateway queue manager, QM1. The gateway queue manager routes the messages to any queue manager in the cluster that hosts the cluster queue Q2.

## Configuring request/reply from a cluster:

Configure a request/reply message path from a cluster to a queue manager outside the cluster. Hide the details of how a queue manager inside the cluster communicates outside the cluster by using a gateway queue manager.

### Before you begin

Figure 136 shows a queue manager, QM2, inside the cluster DEMO. It sends a request to a queue, Q3, hosted on queue manager outside the cluster. The replies are returned to Q2 at QM2 inside the cluster.

To communicate with the queue manager outside the cluster, one or more queue managers inside the cluster act as a gateway. A gateway queue manager has a communication path to the queue managers outside the cluster. In the example, QM1 is the gateway.




Figure 136. Putting to a queue manager outside the cluster

### About this task

Follow the instructions to set up the path for the request and reply messages

### Procedure

1. Send the request message from the cluster.

Consider how the queue manager, QM2, which is inside the cluster puts a message to the queue Q3 at QM3, which is outside the cluster.

- a. Create a QREMOTE definition on QM1 that advertises the remote queue Q3 to the cluster

```
DEFINE QREMOTE(Q3) RNAME(Q3) RQMNAME(QM3) CLUSTER(DEMO)
```

It also has a sender-channel and a transmission queue to the queue manager that is outside the cluster. QM3 has a corresponding receiver-channel. The channels are not shown in Figure 136 on page 980.

An application on QM2 issues an MQPUT call specifying the target queue and the queue to which replies are to be sent. The target queue is Q3 and the reply-to queue is Q2.

The message is sent to QM1, which uses its remote-queue definition to resolve the queue name to Q3 at QM3.

2. Receive the reply message from the queue manager outside the cluster.

A queue manager outside the cluster must have a queue manager alias for each queue manager in the cluster to which it send a message. The queue-manager alias must also specify the name of the transmission queue to the gateway queue manager. In this example, QM3 needs a queue manager alias definition for QM2:

- a. Create a queue manager alias QM2 on QM3

```
DEFINE QREMOTE(QM2) RNAME(' ') RQMNAME(QM2) XMITQ(QM1)
```

QM3 also needs a sender-channel and transmission queue to QM1 and QM1 needs a corresponding receiver-channel.

The application, **app3**, on QM3 can then send replies to QM2, by issuing an MQPUT call and specifying the queue name, Q2 and the queue manager name, QM2.

## What to do next

You can define more than one route out of a cluster.

## Configuring workload balancing from outside a cluster:

Configure a message path from a queue manager outside a cluster to any copy of a cluster queue. The result is to workload balance requests from outside the cluster to each instance of a cluster queue.

## Before you begin

Configure the example, as shown in Figure 134 on page 976 in “Configuring request/reply to a cluster” on page 976.

## About this task

In this scenario, the queue manager outside the cluster, QM3 in Figure 137 on page 982, sends requests to the queue Q2. Q2 is hosted on two queue managers, QM2 and QM4 within cluster DEMO. Both queue managers are configured with a default bind option of NOTFIXED in order to use workload balancing. The requests from QM3, the queue manager outside the cluster, are sent to either instance of Q2 through QM1.

QM3 is not part of a cluster and communicates using distributed queuing techniques. It must have a sender-channel and a transmission queue to QM1. QM1 needs a corresponding receiver-channel. The channels and transmission queues are not shown explicitly in Figure 137 on page 982.

The procedure extends the example in Figure 134 on page 976 in “Configuring request/reply to a cluster” on page 976.

## Procedure

1. Create a QREMOTE definition for Q2 on QM3.  
`DEFINE QREMOTE(Q2) RNAME(Q2) RQMNAME(Q3) XMITQ(QM1)`

Create a QREMOTE definition for each queue in the cluster that QM3 puts messages to.

2. Create a queue-manager alias Q3 on QM1.  
`DEFINE QREMOTE(Q3) RNAME(' ') RQMNAME(' ')`

Q3 is not a real queue manager name. It is the name of a queue manager alias definition in the cluster that equates the queue manager alias name Q3 with blank, ' '.

3. Define a local queue called Q2 on each of QM2 and QM4.  
`DEFINE QLOCAL(Q2) CLUSTER(DEMO) DEFBIND(NOTFIXED)`
4. QM1, the gateway queue manager, has no special definitions.

## Results




Figure 137. Putting from a queue manager outside the cluster

When an application at QM3 issues an MQPUT call to put a message to Q2, the QREMOTE definition on QM3 causes the message to be routed through the gateway queue manager QM1. When QM1 receives the message, it is aware that the message is still intended for a queue named Q2 and performs name resolution. QM1 checks its local definitions and does not find any for Q2. QM1 then checks its cluster configuration and finds that it is aware of two instances of Q2 in cluster DEMO. QM1 can now make use of workload balancing to distribute messages between the instances of Q2 residing on QM2 and QM4.

**Related information:**

Queue name resolution

Name resolution

**Configuring message paths between clusters:**

Connect clusters together using a gateway queue manager. Make queues or queue managers visible to all the clusters by defining cluster queue or cluster queue manager aliases on the gateway queue manager.

**About this task**

Instead of grouping all your queue managers together in one large cluster, you can have many smaller clusters. Each cluster has one or more queue managers in acting as a bridge. The advantage of this is that you can restrict the visibility of queue and queue-manager names across the clusters. See Overlapping clusters. Use aliases to change the names of queues and queue managers to avoid name conflicts or to comply with local naming conventions.




Figure 138. Bridging across clusters

Figure 138 shows two clusters with a bridge between them. There could be more than one bridge.

Configure the clusters using the following procedure:

**Procedure**

1. Define a cluster queue, Q1 on QM1.  
DEFINE QLOCAL(Q1) CLUSTER(CORNISH)
2. Define a cluster queue, Q3 on QM3.  
DEFINE QLOCAL(Q3) CLUSTER(WINDSOR)

3. Create a namelist called CORNISHWINDSOR on QM2, containing the names of both clusters.

```
DEFINE NAMELIST(CORNISHWINDSOR) DESCR('CornishWindsor namelist')
NAMES(CORNISH, WINDSOR)
```

4. Define a cluster queue, Q2 on QM2

```
DEFINE QLOCAL(Q2) CLUSNL(CORNISHWINDSOR)
```

### What to do next

QM2 is a member of both clusters and is the bridge between them. For each queue that you want to make visible across the bridge, you need a QALIAS definition on the bridge. For example in Figure 138 on page 983, on QM2, you need:

```
DEFINE QALIAS(MYQ3) TARGQ(Q3) CLUSTER(CORNISH) DEFBIND(NOTFIXED)
```

Using the queue alias, an application connected to a queue manager in CORNISH, for example QM1, can put a message to Q3. It refers to Q3 as MYQ3. The message is routed to Q3 at QM3.

When you open a queue, you need to set DEFBIND to either NOTFIXED or QDEF. If DEFBIND is left as the default, OPEN, the queue manager resolves the alias definition to the bridge queue manager that hosts it. The bridge does not forward the message.

For each queue manager that you want to make visible, you need a queue-manager alias definition. For example, on QM2 you need:

```
DEFINE QREMOTE(QM1) RNAME(' ') RQMNAME(QM1) CLUSTER(WINDSOR)
```

An application connected to any queue manager in WINDSOR, for example QM3, can put a message to any queue on QM1, by naming QM1 explicitly on the MQOPEN call.

### Queue-manager aliases and clusters:

Use queue-manager aliases to hide the name of queue managers when sending messages into or out of a cluster, and to workload balance messages sent to a cluster.

Queue-manager aliases, which are created using a remote-queue definition with a blank RNAME, have five uses:

#### Remapping the queue-manager name when sending messages

A queue-manager alias can be used to remap the queue-manager name specified in an MQOPEN call to another queue manager. It can be a cluster queue manager. For example, a queue manager might have the queue-manager alias definition:

```
DEFINE QREMOTE(YORK) RNAME(' ') RQMNAME(CLUSQM)
```

YORK can be used as an alias for the queue manager called CLUSQM. When an application on the queue manager that made this definition puts a message to queue manager YORK, the local queue manager resolves the name to CLUSQM. If the local queue manager is not called CLUSQM, it puts the message on the cluster transmission queue to be moved to CLUSQM. It also changes the transmission header to say CLUSQM instead of YORK.

**Note:** The definition applies only on the queue manager that makes it. To advertise the alias to the whole cluster, you need to add the CLUSTER attribute to the remote-queue definition. Then messages from other queue managers that were destined for YORK are sent to CLUSQM.

#### Altering or specifying the transmission queue when sending messages

Aliasing can be used to join a cluster to a non-cluster system. For example, queue managers in the cluster ITALY could communicate with the queue manager called PALERMO, which is outside the cluster. To communicate, one of the queue managers in the cluster must act as a gateway. From the gateway queue manager, issue the command:



```
DEFINE QREMOTE(ROME) RNAME(' ') RQMNAME(PALERMO) XMITQ(X) CLUSTER(ITALY)
```

The command is a queue-manager alias definition. It defines and advertises ROME as a queue manager over which messages from any queue manager in the cluster ITALY can multi-hop to reach their destination at PALERMO. Messages put to a queue opened with the queue-manager name set to ROME are sent to the gateway queue manager with the queue manager alias definition. Once there, the messages are put on the transmission queue X and moved by non-cluster channels to the queue manager PALERMO.

The choice of the name ROME in this example is not significant. The values for QREMOTE and RQMNAME could both be the same.

### **Determining the destination when receiving messages**

When a queue manager receives a message, it extracts the name of the destination queue and queue manager from the transmission header. It looks for a queue-manager alias definition with the same name as the queue manager in the transmission header. If it finds one, it substitutes the RQMNAME from the queue-manager alias definition for the queue manager name in the transmission header.

There are two reasons for using a queue-manager alias in this way:

- To direct messages to another queue manager
- To alter the queue manager name to be the same as the local queue manager

### **Using queue manager aliases in a gateway queue manager to route messages between queue managers in different clusters.**

An application can send a message to a queue in a different cluster using a queue manager alias. The queue does not have to be a cluster queue. The queue is defined in one cluster. The application is connected to a queue manager in a different cluster. A gateway queue manager connects the two clusters. If the queue is not defined as clustered, for the correct routing to take place, the application must open the queue using the queue name and a clustered queue manager alias name. For an example of a configuration, see “Creating two-overlapping clusters with a gateway queue manager” on page 947, from which the reply message flow illustrated in figure 1, is taken.

The diagram shows the path taken by the reply message back to a temporary dynamic queue, which is called RQ. The server application, connected to QM3, opens the reply queue using the queue manager name QM2. The queue manager name QM2 is defined as a clustered queue manager alias on QM1. QM3 routes the reply message to QM1. QM1 routes the message to QM2.

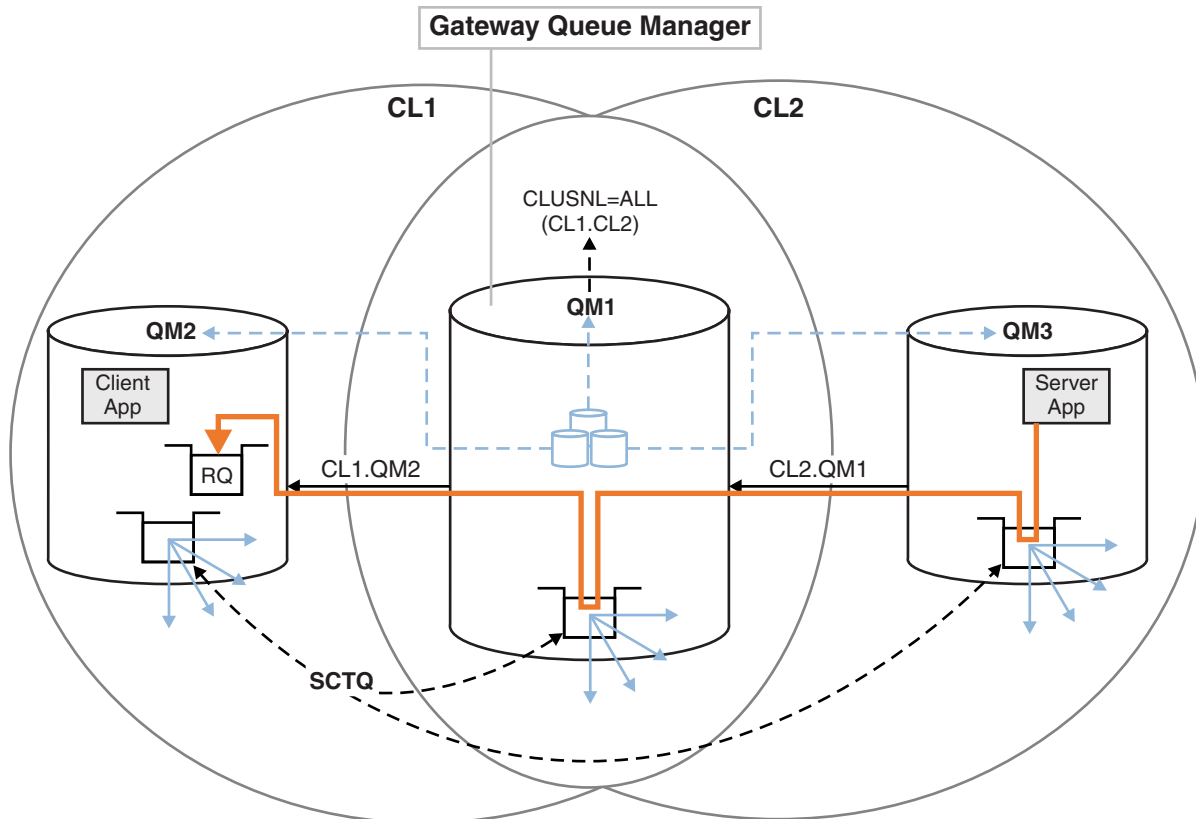


Figure 139. Using a queue manager alias to return the reply message to a different cluster

The way the routing works is as follows. Every queue manager in each cluster has a queue manager alias definition on QM1. The aliases are clustered in all the clusters. The grey dashed arrows from each of the aliases to a queue manager show that each queue manager alias is resolved to a real queue manager in at least one of the clusters. In this case, the QM2 alias is clustered in both cluster CL1 and CL2, and is resolved to the real queue manager QM2 in CL1. The server application creates the reply message using the reply to queue name RQ, and reply to queue manager name QM2. The message is routed to QM1 because the queue manager alias definition QM2 is defined on QM1 in cluster CL2 and queue manager QM2 is not in cluster CL2. As the message cannot be sent to the target queue manager, it is sent to the queue manager that has the alias definition.

QM1 places the message on the cluster transmission queue on QM1 for transferal to QM2. QM1 routes the message to QM2 because the queue manager alias definition on QM1 for QM2 defines QM2 as the real target queue manager. The definition is not circular, because alias definitions can refer only to real definitions; the alias cannot point to itself. The real definition is resolved by QM1, because both QM1 and QM2 are in the same cluster, CL1. QM1 finds out the connection information for QM2 from the repository for CL1, and routes the message to QM2. For the message to be rerouted by QM1, the server application must have opened the reply queue with the option `DEFBIND` set to `MQBND_BIND_NOT_FIXED`. If the server application had opened the reply queue with the option `MQBND_BIND_ON_OPEN`, the message is not rerouted and ends up on a dead letter queue.

### Using a queue manager as a gateway into the cluster to workload balance messages from coming from outside the cluster.

You define a queue called EDINBURGH on more than one queue manager in the cluster. You want the clustering mechanism to balance the workload for messages coming to that queue from outside the cluster.

A queue manager from outside the cluster needs a transmit queue and sender-channel to one queue manager in the cluster. This queue is called a gateway queue manager. To take advantage of the default workload balancing mechanism, one of the following rules must apply:

- The gateway queue manager must not contain an instance of the EDINBURGH queue.
- The gateway queue manager specifies CLWLUSEQ(ANY) on ALTER QMGR.

For an example of workload balancing from outside a cluster, see “Configuring workload balancing from outside a cluster” on page 981

### Reply-to queue aliases and clusters:

A reply-to queue alias definition is used to specify alternative names for reply information. Reply-to queue alias definitions can be used with clusters just the same as in a distributed queuing environment.

For example:

- An application at queue manager VENICE sends a message to queue manager PISA using the MQPUT call. The application provides the following reply-to queue information in the message descriptor:  

```
ReplyToQ='QUEUE'
ReplyToQMGr=''
```
- In order that replies sent to QUEUE can be received on OTHERQ at PISA, create a remote-queue definition on VENICE that is used as a reply-to queue alias. The alias is effective only on the system on which it was created.  

```
DEFINE QREMOTE(QUEUE) RNAME(OTHERQ) RQMNAME(PISA)
```

RQMNAME and QREMOTE can specify the same names, even if RQMNAME is itself a cluster queue manager.

### Queue aliases and clusters:

Use queue aliases to hide the name of a cluster queue, to cluster a queue, adopt different attributes, or adopt different access controls.

A QALIAS definition is used to create an alias by which a queue is to be known. You might create an alias for a number of reasons:

- You want to start using a different queue but you do not want to change your applications.
- You do not want applications to know the real name of the queue to which they are putting messages.
- You might have a naming convention that differs from the one where the queue is defined.
- Your applications might not be authorized to access the queue by its real name but only by its alias.

Create a QALIAS definition on a queue manager using the DEFINE QALIAS command. For example, run the command:

```
DEFINE QALIAS(PUBLIC) TARGQ(LOCAL) CLUSTER(C)
```

The command advertises a queue called PUBLIC to the queue managers in cluster C. PUBLIC is an alias that resolves to the queue called LOCAL. Messages sent to PUBLIC are routed to the queue called LOCAL.

You can also use a queue alias definition to resolve a queue name to a cluster queue. For example, run the command:


```
DEFINE QALIAS(PRIVATE) TARGQ(PUBLIC)
```

The command enables a queue manager to use the name PRIVATE to access a queue advertised elsewhere in the cluster by the name PUBLIC. Because this definition does not include the CLUSTER attribute it applies only to the queue manager that makes it.

## Using clusters for workload management

By defining multiple instances of a queue on different queue managers in a cluster you can spread the work of servicing the queue over multiple servers. There are several factors that can prevent messages being requested to a different queue manager in the event of failure.


As well as setting up clusters to reduce system administration, you can create clusters in which more than one queue manager hosts an instance of the same queue.

You can organize your cluster such that the queue managers in it are clones of each other. Each queue manager is able to run the same applications and have local definitions of the same queues.  For example, in a z/OS parallel sysplex the cloned applications might access data in a shared Db2 or Virtual Storage Access Method (VSAM) database. You can spread the workload between your queue managers by having several instances of an application. Each instance of the application receives messages and runs independently of the others.

The advantages of using clusters in this way are as follows:

- Increased availability of your queues and applications.
- Faster throughput of messages.
- More even distribution of workload in your network.

Any one of the queue managers that hosts an instance of a particular queue can handle messages destined for that queue, and applications do not name a queue manager when sending messages. If a cluster contains more than one instance of the same queue, IBM MQ selects a queue manager to route a message to. Suitable destinations are chosen based on the availability of the queue manager and queue, and on a number of cluster workload-specific attributes associated with queue managers, queues, and channels. See [Workload balancing in clusters](#).

 In IBM MQ for z/OS, queue managers that are in queue-sharing groups can host cluster queues as shared queues. Shared cluster queues are available to all queue managers in the same queue-sharing group. For example, in a cluster with multiple instances of the same queue, either or both of the queue managers QM2 and QM4 can be a shared-queue manager. Each has a definition for the queue Q3. Any of the queue managers in the same queue-sharing group as QM4 can read a message put to the shared queue Q3. Each queue-sharing group can contain up to 32 queue managers, each with access to the same data. Queue sharing significantly increases the throughput of your messages.

See the following subtopics for more information about cluster configurations for workload management:

### Related concepts:

[Comparison of clustering and distributed queuing](#)

[Distributed queuing and clusters](#)

[Components of a cluster](#)

[Cluster channels](#)

[“Routing messages to and from clusters” on page 975](#)

Use queue aliases, queue manager aliases, and remote queue definitions to connect clusters to external queue managers and other clusters.

### Related tasks:

[“Configuring a queue manager cluster” on page 902](#)

Clusters provide a mechanism for interconnecting queue managers in a way that simplifies both the initial configuration and the ongoing management. You can define cluster components, and create and manage clusters.

[“Setting up a new cluster” on page 913](#)

Follow these instructions to set up the example cluster. Separate instructions describe setting up the cluster on TCP/IP, LU 6.2, and with a single transmission queue or multiple transmission queues. Test

the cluster works by sending a message from one queue manager to the other.

“Configuring workload balancing from outside a cluster” on page 981

Configure a message path from a queue manager outside a cluster to any copy of a cluster queue. The result is to workload balance requests from outside the cluster to each instance of a cluster queue.

**Related reference:**

What happens if a cluster queue is disabled for MQPUT

Workload balancing set on a cluster-sender channel is not working

**Related information:**

The Cluster Queue Monitoring sample program (AMQSCLM)

Writing and compiling cluster workload exits

**Example of a cluster with more than one instance of a queue:**

In this example of a cluster with more than one instance of a queue, messages are routed to different instances of the queue. You can force a message to a specific instance of the queue, and you can choose to send a sequence of messages to one of either of the queue managers.

Figure 140 on page 990 shows a cluster in which there is more than one definition for the queue Q3. If an application at QM1 puts a message to Q3, it does not necessarily know which instance of Q3 is going to process its message. If an application is running on QM2 or QM4, where there are local instances of Q3, the local instance of Q3 is opened by default. By setting the CLWLUSEQ queue attribute, the local instance of the queue can be treated the same as a remote instance of the queue.

The MQOPEN option DefBind controls whether the target queue manager is chosen when the MQOPEN call is issued, or when the message is transferred from the transmission queue.

If you set DefBind to MQBND\_BIND\_NOT\_FIXED the message can be sent to an instance of the queue that is available when the message is transmitted. This avoids the following problems:

- The target queue is unavailable when the message arrives at the target queue manager.
- The state of the queue has changed.
- The message has been put using a cluster queue alias, and no instance of the target queue exists on the queue manager where the instance of the cluster queue alias is defined.

If any of these problems are discovered at transmission time, another available instance of the target queue is sought and the message is rerouted. If no instances of the queue are available, the message is placed on the dead-letter queue.




Figure 140. A cluster with multiple instances of the same queue

One factor that can prevent messages being rerouted is if messages have been assigned to a fixed queue manager or channel with `MQBND_BIND_ON_OPEN`. Messages bound on `MQOPEN` are never reallocated to another channel. Note also that message reallocation only takes place when a cluster channel is actually failing. Reallocation does not occur if the channel has already failed.

The system attempts to reroute a message if the destination queue manager goes out of service. In so doing, it does not affect the integrity of the message by running the risk of losing it or by creating a duplicate. If a queue manager fails and leaves a message in doubt, that message is not rerouted.

**z/OS** On IBM MQ for z/OS, the channel does not completely stop until the message reallocation process is complete. Stopping the channel with mode set to `FORCE` or `TERMINATE` does interrupt the process, so if you do this then some `BIND_NOT_FIXED` messages might have already been reallocated to another channel, or the messages might be out of order.

**Note:** **z/OS**

1. Before setting up a cluster that has multiple instances of the same queue, ensure that your messages do not have dependencies on each other. For example, needing to be processed in a specific sequence or by the same queue manager.
2. Make the definitions for different instances of the same queue identical. Otherwise you get different results from different `MQINQ` calls.

## Adding a queue manager that hosts a queue locally:

Follow these instructions to add an instance of INVENTQ to provide additional capacity to run the inventory application system in Paris and New York.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in Adding a new queue manager to a cluster. It contains three queue managers; LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository. The inventory application runs on the system in New York, connected to the NEWYORK queue manager. The application is driven by the arrival of messages on the INVENTQ queue.
- We want to add an instance of INVENTQ to provide additional capacity to run the inventory application system in Paris and New York.

### About this task

Follow these steps to add a queue manager that hosts a queue locally.

### Procedure

1. Alter the PARIS queue manager.

For the application in Paris to use the INVENTQ in Paris and the one in New York, we must inform the queue manager. On PARIS issue the following command:

```
ALTER QMGR CLWLUSEQ(ANY)
```

2. Review the inventory application for message affinities.

Before proceeding, ensure that the inventory application does not have any dependencies on the sequence of processing of messages. For more information, see Handling message affinities.

3. Install the inventory application on the system in Paris.

4. Define the cluster queue INVENTQ.

The INVENTQ queue which is already hosted by the NEWYORK queue manager is also to be hosted by PARIS. Define it on the PARIS queue manager as follows:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
```

Now that you have completed all the definitions, if you have not already done so, start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program on queue manager PARIS. The listener listens for incoming network requests and starts the cluster-receiver channel when it is needed.

### Results

Figure 141 on page 992 shows the cluster set up by this task.

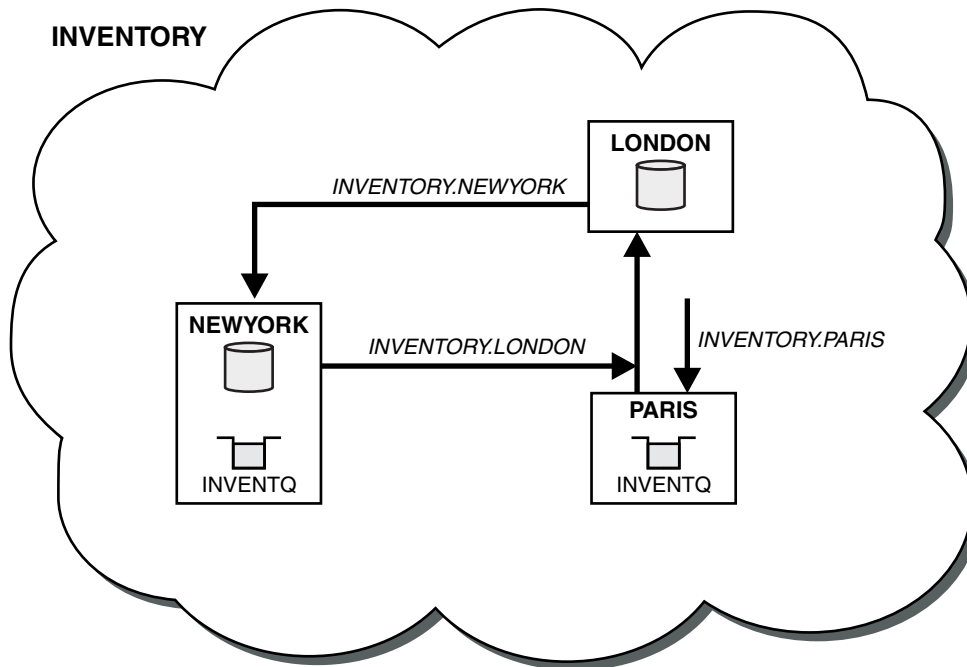


Figure 141. The INVENTORY cluster, with three queue managers

The modification to this cluster was accomplished without you altering the queue managers NEWYORK or LONDON. The full repositories in these queue managers are updated automatically with the information they need to be able to send messages to INVENTQ at PARIS.

### What to do next

The INVENTQ queue and the inventory application are now hosted on two queue managers in the cluster. This increases their availability, speeds up throughput of messages, and allows the workload to be distributed between the two queue managers. Messages put to INVENTQ by any of the queue managers LONDON, NEWYORK, PARIS are routed alternately to PARIS or NEWYORK, so that the workload is balanced.

### Using two networks in a cluster:

Follow these instructions to add a new store in TOKYO where there are two different networks. Both need to be available for use to communicate with the queue manager in Tokyo.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in "Adding a queue manager to a cluster". It contains three queue managers; LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository. The inventory application runs on the system in New York, connected to the NEWYORK queue manager. The application is driven by the arrival of messages on the INVENTQ queue.
- A new store is being added in TOKYO where there are two different networks. Both need to be available for use to communicate with the queue manager in Tokyo.



## About this task

Follow these steps to use two networks in a cluster.

### Procedure

1. Decide which full repository TOKYO refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories to gather information about the cluster. It builds up its own partial repository. It is of no particular significance which repository you choose. In this example, NEWYORK is chosen. Once the new queue manager has joined the cluster it communicates with both of the repositories.

2. Define the CLUSRCVR channels.

Every queue manager in a cluster needs to define a cluster-receiver on which it can receive messages. This queue manager needs to be able to communicate on each network.

```
DEFINE CHANNEL(INVENTORY.TOKYO.NETB) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME('TOKYO.NETB.CMSTORE.COM') CLUSTER(INVENTORY) DESCR('Cluster-receiver channel using
network B for TOKYO')
```

```
DEFINE CHANNEL(INVENTORY.TOKYO.NETA) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME('TOKYO.NETA.CMSTORE.COM') CLUSTER(INVENTORY) DESCR('Cluster-receiver channel using
network A for TOKYO')
```

3. Define a CLUSSDR channel on queue manager TOKYO.

Every queue manager in a cluster needs to define one cluster-sender channel on which it can send messages to its first full repository. In this case we have chosen NEWYORK, so TOKYO needs the following definition:

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP) CONNAME(NEWYORK.CHSTORE.COM)
CLUSTER(INVENTORY) DESCR('Cluster-sender channel from TOKYO to repository at NEWYORK')
```

Now that you have completed all the definitions, if you have not already done so start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program on queue manager PARIS. The listener program listens for incoming network requests and starts the cluster-receiver channel when it is needed.

### Results

Figure 142 on page 994 shows the cluster set up by this task.

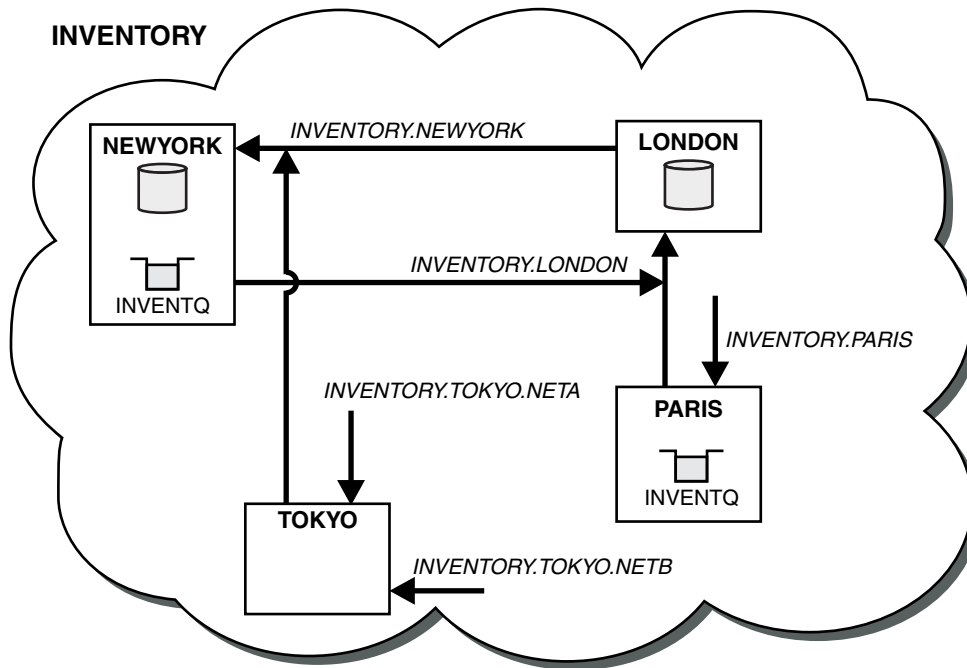


Figure 142. The *INVENTORY* cluster, with four queue managers

By making only three definitions, we have added the queue manager *TOKYO* to the cluster with two different network routes available to it.

**Related tasks:**

“Adding a queue manager to a cluster” on page 924

Follow these instructions to add a queue manager to the cluster you created. Messages to cluster queues and topics are transferred using the single cluster transmission queue `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

**Using a primary and a secondary network in a cluster:**

Follow these instructions to make one network the primary network, and another network the backup network. Use the backup network if there is a problem with the primary network.

**Before you begin**

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The *INVENTORY* cluster has been set up as described in “Using two networks in a cluster” on page 992. It contains four queue managers; *LONDON* and *NEWYORK* both hold full repositories; *PARIS* and *TOKYO* hold partial repositories. The inventory application runs on the system in New York, connected to the queue manager *NEWYORK*. The *TOKYO* queue manager has two different networks that it can communicate on.
- You want to make one of the networks the primary network, and another of the networks the backup network. You plan to use the backup network if there is a problem with the primary network.

**About this task**

Use the `NETPRTY` attribute to configure a primary and a secondary network in a cluster.

## Procedure

Alter the existing CLUSRCVR channels on TOKYO.

To indicate that the network A channel is the primary channel, and the network B channel is the secondary channel, use the following commands:

1. ALTER CHANNEL(INVENTORY.TOKYO.NETA) CHLTYPE(CLUSRCVR) NETPRTY(2) DESCR('Main cluster-receiver channel for TOKYO')
2. ALTER CHANNEL(INVENTORY.TOKYO.NETB) CHLTYPE(CLUSRCVR) NETPRTY(1) DESCR('Backup cluster-receiver channel for TOKYO')

## What to do next

By configuring the channel with different network priorities, you have now defined to the cluster that you have a primary network and a secondary network. The queue managers in the cluster that use these channels automatically use the primary network whenever it is available. The queue managers failover to use the secondary network when the primary network is not available.

## Adding a queue to act as a backup:

Follow these instructions to provide a backup in Chicago for the inventory system that now runs in New York. The Chicago system is only used when there is a problem with the New York system.

## Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager to a cluster” on page 924. It contains three queue managers; LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository. The inventory application runs on the system in New York, connected to the NEWYORK queue manager. The application is driven by the arrival of messages on the INVENTQ queue.
- A new store is being set up in Chicago to provide a backup for the inventory system that now runs in New York. The Chicago system only used when there is a problem with the New York system.

## About this task

Follow these steps to add a queue to act as a backup.

## Procedure

1. Decide which full repository CHICAGO refers to first.

Every queue manager in a cluster must refer to one or other of the full repositories to gather information about the cluster. It builds up its own partial repository. It is of no particular significance which repository you choose for any particular queue manager. In this example, NEWYORK is chosen. Once the new queue manager has joined the cluster it communicates with both of the repositories.

2. Define the CLUSRCVR channel.

Every queue manager in a cluster needs to define a cluster-receiver on which it can receive messages. On CHICAGO, define:

```
DEFINE CHANNEL(INVENTORY.CHICAGO) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(CHICAGO.CMSTORE.COM) CLUSTER(INVENTORY) DESCR('Cluster-receiver channel for
CHICAGO')
```

3. Define a CLUSSDR channel on queue manager CHICAGO.

Every queue manager in a cluster needs to define one cluster-sender channel on which it can send messages to its first full repository. In this case we have chosen NEWYORK, so CHICAGO needs the following definition:

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP) CONNAME(NEWYORK.CHSTORE.COM)
CLUSTER(INVENTORY) DESCR('Cluster-sender channel from CHICAGO to repository at NEWYORK')
```

4. Alter the existing cluster queue INVENTQ.

The INVENTQ which is already hosted by the NEWYORK queue manager is the main instance of the queue.

```
ALTER QLOCAL(INVENTQ) CLWLPRTY(2)
```

5. Review the inventory application for message affinities.

Before proceeding, ensure that the inventory application does not have any dependencies on the sequence of processing of messages.

6. Install the inventory application on the system in CHICAGO.

7. Define the backup cluster queue INVENTQ

The INVENTQ which is already hosted by the NEWYORK queue manager, is also to be hosted as a backup by CHICAGO. Define it on the CHICAGO queue manager as follows:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY) CLWLPRTY(1)
```

Now that you have completed all the definitions, if you have not already done so start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program on queue manager CHICAGO. The listener program listens for incoming network requests and starts the cluster-receiver channel when it is needed.

## Results

Figure 143 shows the cluster set up by this task.

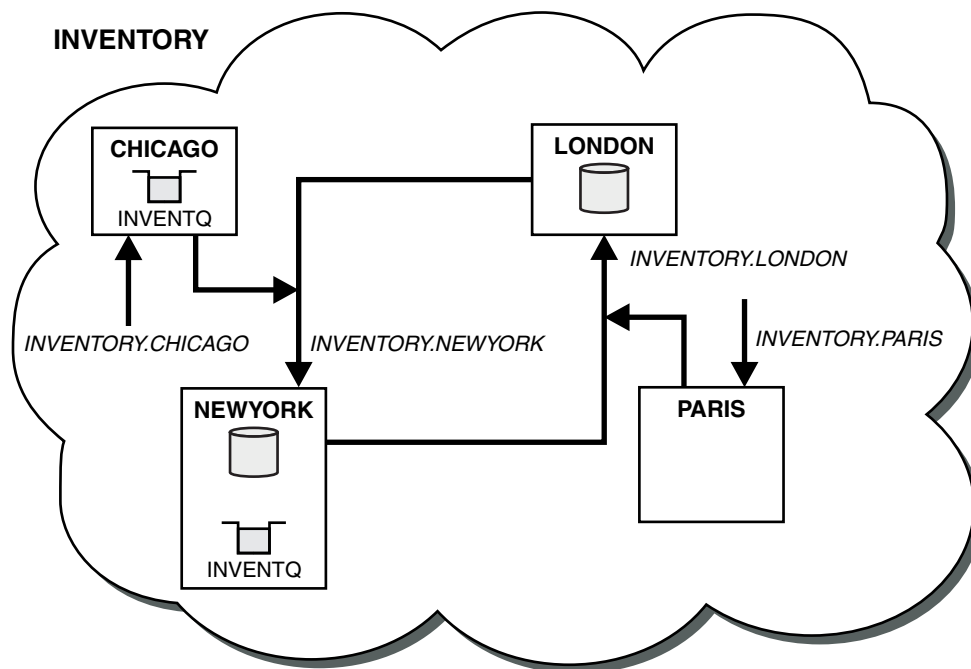


Figure 143. The INVENTORY cluster, with four queue managers

The INVENTQ queue and the inventory application are now hosted on two queue managers in the cluster. The CHICAGO queue manager is a backup. Messages put to INVENTQ are routed to NEWYORK unless it is unavailable when they are sent instead to CHICAGO.

**Note:**

The availability of a remote queue manager is based on the status of the channel to that queue manager. When channels start, their state changes several times, with some of the states being less preferential to the cluster workload management algorithm. In practice this means that lower priority (backup) destinations can be chosen while the channels to higher priority (primary) destinations are starting.

If you need to ensure that no messages go to a backup destination, do not use CLWLPRTY. Consider using separate queues, or CLWLRANK with a manual switch over from the primary to backup.

**Restricting the number of channels used:**

Follow these instructions to restrict the number of active channels each server runs when a price check application is installed on various queue managers.

**Before you begin**

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- A price check application is to be installed on various queue managers. To keep the number of channels being used to a low number, the number of active channels each server runs is restricted. The application is driven by the arrival of messages on the PRICEQ queue.
- Four server queue managers host the price check application. Two query queue managers send messages to the PRICEQ to query a price. Two more queue managers are configured as full repositories.

**About this task**

Follow these steps to restrict the number of channels used.

**Procedure**

1. Choose two full repositories.

Choose two queue managers to be the full repositories for your price check cluster. They are called REPOS1 and REPOS2.

Issue the following command:

```
ALTER QMGR REPOS(PRICECHECK)
```

2. Define a CLUSRCVR channel on each queue manager.

At each queue manager in the cluster, define a cluster-receiver channel and a cluster-sender channel. It does not matter which is defined first.

```
DEFINE CHANNEL(PRICECHECK.SERVE1) CHLTYPE(CLUSRCVR) TRPTYPE(TCP) CONNAME(SERVER1.COM)
CLUSTER(PRICECHECK) DESCR('Cluster-receiver channel')
```

3. Define a CLUSSDR channel on each queue manager.

Make a CLUSSDR definition at each queue manager to link that queue manager to one or other of the full repository queue managers.

```
DEFINE CHANNEL(PRICECHECK.REPOS1) CHLTYPE(CLUSSDR) TRPTYPE(TCP) CONNAME(REPOS1.COM)
CLUSTER(PRICECHECK) DESCR('Cluster-sender channel to repository queue manager')
```

4. Install the price check application.
5. Define the PRICEQ queue on all the server queue managers.

Issue the following command on each:  
DEFINE QLOCAL(PRICEQ) CLUSTER(PRICECHECK)

6. Restrict the number of channels used by queries

On the query queue managers we restrict the number of active channels used, by issuing the following commands on each:

```
ALTER QMGR CLWLMRUC(2)
```

7. If you have not already done so, start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program.

The listener program listens for incoming network requests and starts the cluster-receiver channel when it is needed.

## Results

Figure 144 shows the cluster set up by this task.

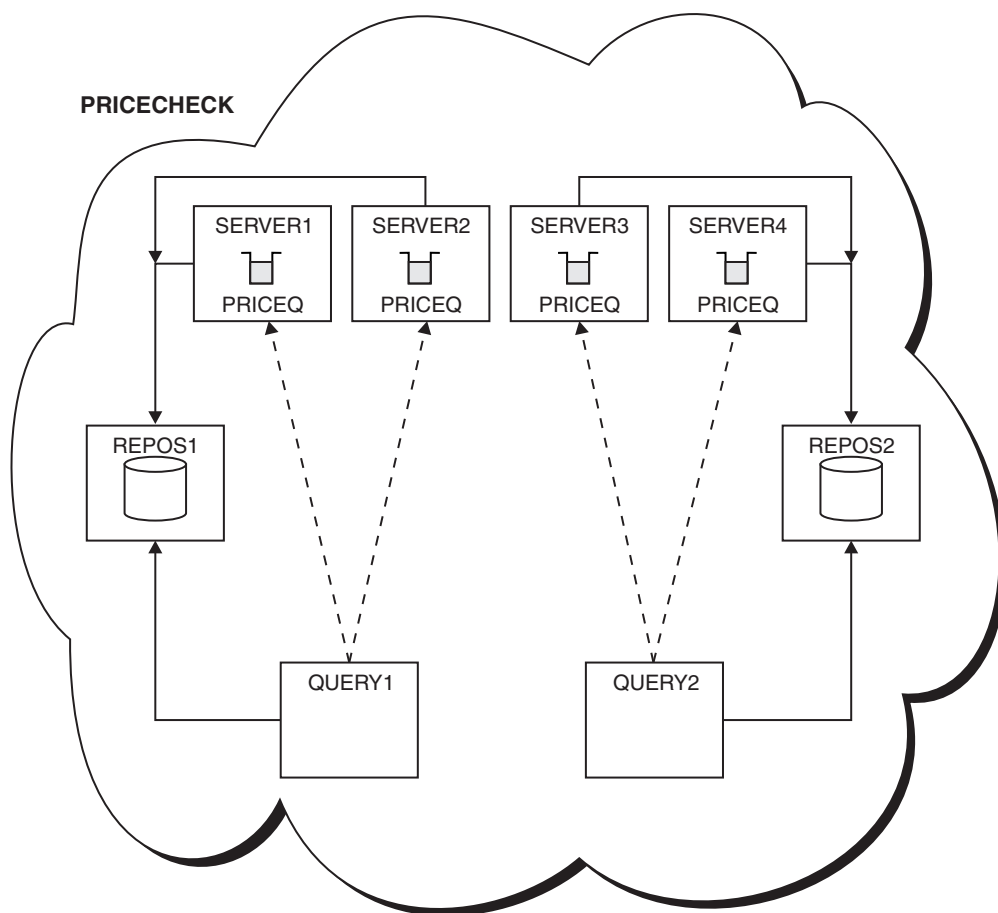


Figure 144. The PRICECHECK cluster, with four server queue managers, two repositories, and two query queue managers

Although there are four instances of the PRICEQ queue available in the PRICECHECK cluster, each querying queue manager only uses two of two of them. For example, the QUERY1 queue manager only has active channels to the SERVER1 and SERVER2 queue managers. If SERVER1 became unavailable, the QUERY1 queue manager would then begin to use another queue manager, for example SERVER3.

## Adding a more powerful queue manager that hosts a queue:

Follow these instructions to provide additional capacity by running the inventory system in Los Angeles as well as New York, where Los Angeles can handle twice the number of messages as New York.

### Before you begin

**Note:** For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager to a cluster” on page 924. It contains three queue managers: LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository and puts messages from INVENTQ. The inventory application runs on the system in New York connected to the NEWYORK queue manager. The application is driven by the arrival of messages on the INVENTQ queue.
- A new store is being set up in Los Angeles. To provide additional capacity, you want to run the inventory system in Los Angeles as well as New York. The new queue manager can process twice as many messages as New York.

### About this task

Follow these steps to add a more powerful queue manager that hosts a queue.

### Procedure

1. Decide which full repository LOSANGELES refers to first.
2. Every queue manager in a cluster must refer to one or other of the full repositories to gather information about the cluster. It builds up its own partial repository. It is of no particular significance which repository you choose. In this example, NEWYORK is chosen. Once the new queue manager has joined the cluster it communicates with both of the repositories.

```
DEFINE CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(NEWYORK.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-sender channel from LOSANGELES to repository at NEWYORK')
```

3. Define the CLUSRCVR channel on queue manager LOSANGELES. Every queue manager in a cluster must define a cluster-receiver channel on which it can receive messages. On LOSANGELES, define:

```
DEFINE CHANNEL(INVENTORY.LOSANGELES) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(LOSANGELES.CHSTORE.COM) CLUSTER(INVENTORY)
DESCR('Cluster-receiver channel for queue manager LOSANGELES')
CLWLWGT(2)
```

The cluster-receiver channel advertises the availability of the queue manager to receive messages from other queue managers in the cluster INVENTORY. Setting CLWLWGT to two ensures that the Los Angeles queue manager gets twice as many of the inventory messages as New York (when the channel for NEWYORK is set to one).

4. Alter the CLUSRCVR channel on queue manager NEWYORK.  
Ensure that the Los Angeles queue manager gets twice as many of the inventory messages as New York. Alter the definition of the cluster-receiver channel.  

```
ALTER CHANNEL(INVENTORY.NEWYORK) CHLTYPE(CLUSRCVR) CLWLWGT(1)
```
5. Review the inventory application for message affinities.  
Before proceeding, ensure that the inventory application does not have any dependencies on the sequence of processing of messages.
6. Install the inventory application on the system in Los Angeles
7. Define the cluster queue INVENTQ.

The INVENTQ queue, which is already hosted by the NEWYORK queue manager, is also to be hosted by LOSANGELES. Define it on the LOSANGELES queue manager as follows:

```
DEFINE QLOCAL(INVENTQ) CLUSTER(INVENTORY)
```

Now that you have completed all the definitions, if you have not already done so start the channel initiator on IBM MQ for z/OS. On all platforms, start a listener program on queue manager LOSANGELES. The listener program listens for incoming network requests and starts the cluster-receiver channel when it is needed.

## Results

“Adding a more powerful queue manager that hosts a queue” on page 999 shows the cluster set up by this task.

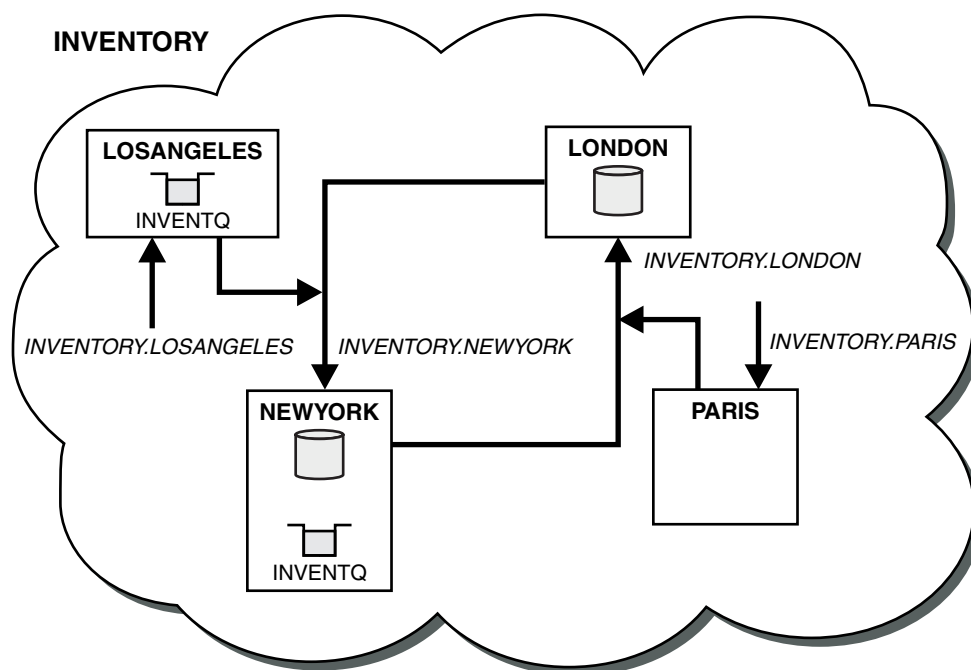


Figure 145. The INVENTORY cluster with four queue managers

This modification to the cluster was accomplished without you having to alter the queue managers LONDON and PARIS. The repositories in these queue managers are updated automatically with the information they need to be able to send messages to INVENTQ at LOSANGELES.

## What to do next

The INVENTQ queue and inventory application are hosted on two queue managers in the cluster. The configuration increases their availability, speeds up throughput of messages, and allows the workload to be distributed between the two queue managers. Messages put to INVENTQ by either LOSANGELES or NEWYORK are handled by the instance on the local queue manager whenever possible. Messages put by LONDON or PARIS are routed to LOSANGELES or NEWYORK, with twice as many messages being sent to LOSANGELES.



## **Application programming and clusters:**

You do not need to make any programming changes to take advantage of multiple instances of the same queue. However, some programs do not work correctly unless a sequence of messages is sent to the same instance of a queue.

Applications can open a queue using the MQOPEN call. Applications use the MQPUT call to put messages onto an open queue. Applications can put a single message onto a queue that is not already open, using the MQPUT1 call.

If you set up clusters that have multiple instances of the same queue, there are no specific application programming considerations. However, to benefit from the workload management aspects of clustering, you might need to modify your applications. If you set up a network in which there are multiple definitions of the same queue, review your applications for message affinities.

Suppose for example, you have two applications that rely on a series of messages flowing between them in the form of questions and answers. You probably want answers to go back to the same queue manager that sent a question. It is important that the workload management routine does not send the messages to any queue manager that hosts a copy of the reply queue.

You might have applications that require messages to be processed in sequence (for example, a database replication application that sends batches of messages that must be retrieved in sequence). The use of segmented messages can also cause an affinity problem.

### **Opening a local or remote version of the target queue**

Be aware of how the queue manager chooses whether use a local or remote version of the target queue.

1. The queue manager opens the local version of the target queue to read messages, or to set the attributes of the queue.
2. The queue manager opens any instance of the target queue to write messages to, if at least one of the following conditions is true:
  - A local version of the target queue does not exist.
  - The queue manager specifies CLWLUSEQ(ANY) on ALTER QMGR.
  - The queue on the queue manager specifies CLWLUSEQ(ANY).

### *Handling message affinities:*

Message affinities are rarely part of good programming design. You need to remove message affinities to use clustering fully. If you cannot remove message affinities, you can force related messages to be delivered using the same channel and to the same queue manager.

If you have applications with message affinities, remove the affinities before starting to use clusters.

Removing message affinities improves the availability of applications. An application sends a batch of messages that has message affinities to a queue manager. The queue manager fails after receiving only part of the batch. The sending queue manager must wait for it to recover and process the incomplete message batch before it can send any more messages.

Removing messages affinities also improves the scalability of applications. A batch of messages with affinities can lock resources at the destination queue manager while waiting for subsequent messages. These resources might remain locked for long periods of time, preventing other applications from doing their work.

Furthermore, message affinities prevent the cluster workload management routines from making the best choice of queue manager.

To remove affinities, consider the following possibilities:

- Carrying state information in the messages
- Maintaining state information in nonvolatile storage accessible to any queue manager, for example in a Db2 database
- Replicating read-only data so that it is accessible to more than one queue manager

If it is not appropriate to modify your applications to remove message affinities, there are a number of possible solutions to the problem.

### **Name a specific destination on the MQOPEN call**

Specify the remote-queue name and the queue manager name on each MQOPEN call, and all messages put to the queue using that object handle go to the same queue manager, which might be the local queue manager.

Specifying the remote-queue name and the queue manager name on each MQOPEN call has disadvantages:

- No workload balancing is carried out. You do not take advantage of the benefits of cluster workload balancing.
- If the target queue manager is remote and there is more than one channel to it, the messages might take different routes and the sequence of messages is still not preserved.
- If your queue manager has a definition for a transmission queue with the same name as the destination queue manager, messages go on that transmission queue rather than on the cluster transmission queue.

### **Return the queue-manager name in the reply-to queue manager field**

Allow the queue manager that receives the first message in a batch to return its name in its response. It does this using the ReplyToQMgr field of the message descriptor. The queue manager at the sending end can then extract the reply-to queue manager name and specify it on all subsequent messages.

Using the ReplyToQMgr information from the response has disadvantages:

- The requesting queue manager must wait for a response to its first message
- You must write additional code to find and use the ReplyToQMgr information before sending subsequent messages
- If there is more than one route to the queue manager, the sequence of the messages might not be preserved

### **Set the MQ00\_BIND\_ON\_OPEN option on the MQOPEN call**

Force all your messages to be put to the same destination using the MQ00\_BIND\_ON\_OPEN option on the MQOPEN call. Either MQ00\_BIND\_ON\_OPEN or MQ00\_BIND\_ON\_GROUP must be specified when using message groups with clusters to ensure that all messages in the group are processed at the same destination.

By opening a queue and specifying MQ00\_BIND\_ON\_OPEN, you force all messages that are sent to this queue to be sent to the same instance of the queue. MQ00\_BIND\_ON\_OPEN binds all messages to the same queue manager and also to the same route. For example, if there is an IP route and a NetBIOS route to the same destination, one of these is selected when the queue is opened and this selection is honored for all messages put to the same queue using the object handle obtained.

By specifying MQ00\_BIND\_ON\_OPEN you force all messages to be routed to the same destination. Therefore applications with message affinities are not disrupted. If the destination is not available, the messages remain on the transmission queue until it becomes available again.

MQOO\_BIND\_ON\_OPEN also applies when the queue manager name is specified in the object descriptor when you open a queue. There might be more than one route to the named queue manager. For example, there might be multiple network paths or another queue manager might have defined an alias. If you specify MQOO\_BIND\_ON\_OPEN, a route is selected when the queue is opened.

**Note:** This is the recommended technique. However, it does not work in a multi-hop configuration in which a queue manager advertises an alias for a cluster queue. Nor does it help in situations in which applications use different queues on the same queue manager for different groups of messages.

An alternative to specifying MQOO\_BIND\_ON\_OPEN on the MQOPEN call, is to modify your queue definitions. On your queue definitions, specify DEFBIND(OPEN), and allow the DefBind option on the MQOPEN call to default to MQOO\_BIND\_AS\_Q\_DEF.

### **Set the MQOO\_BIND\_ON\_GROUP option on the MQOPEN call**

Force all your messages in a group to be put to the same destination using the MQOO\_BIND\_ON\_GROUP option on the MQOPEN call. Either MQOO\_BIND\_ON\_OPEN or MQOO\_BIND\_ON\_GROUP must be specified when using message groups with clusters to ensure that all messages in the group are processed at the same destination.

By opening a queue and specifying MQOO\_BIND\_ON\_GROUP, you force all messages in a group that are sent to this queue to be sent to the same instance of the queue. MQOO\_BIND\_ON\_GROUP binds all messages in a group to the same queue manager, and also to the same route. For example, if there is an IP route and a NetBIOS route to the same destination, one of these is selected when the queue is opened and this selection is honored for all messages in a group put to the same queue using the object handle obtained.

By specifying MQOO\_BIND\_ON\_GROUP you force all messages in a group to be routed to the same destination. Therefore applications with message affinities are not disrupted. If the destination is not available, the messages remain on the transmission queue until it becomes available again.

MQOO\_BIND\_ON\_GROUP also applies when the queue manager name is specified in the object descriptor when you open a queue. There might be more than one route to the named queue manager. For example, there might be multiple network paths or another queue manager might have defined an alias. If you specify MQOO\_BIND\_ON\_GROUP, a route is selected when the queue is opened.

If MQOO\_BIND\_ON\_GROUP is specified but the messages are not grouped, the behavior is equivalent to MQOO\_BIND\_NOT\_FIXED.

**Note:** This is the recommended technique for ensuring that messages in a group are sent to the same destination. However, it does not work in a multi-hop configuration in which a queue manager advertises an alias for a cluster queue.

An alternative to specifying MQOO\_BIND\_ON\_GROUP on the MQOPEN call, is to modify your queue definitions. On your queue definitions, specify DEFBIND(GROUP), and allow the DefBind option on the MQOPEN call to default to MQOO\_BIND\_AS\_Q\_DEF.

### **Write a customized cluster workload exit program**

Instead of modifying your applications you can circumvent the message affinities problem by writing a cluster workload exit program. Writing a cluster workload exit program is not easy and is not a recommended solution. The program would have to be designed to recognize the affinity by inspecting the content of messages. Having recognized the affinity, the program would have to force the workload management utility to route all related messages to the same queue manager.

---

## Configuring publish/subscribe messaging

You can start, stop and display the status of queued publish/subscribe. You can also add and remove streams, and add and delete queue managers from a broker hierarchy.

See the following subtopics for more information on controlling queued publish/subscribe:

### Setting queued publish/subscribe message attributes

You control the behavior of some publish/subscribe message attributes using queue manager attributes. The other attributes you control in the *Broker* stanza of the *qm.ini* file.

### About this task

You can set the following publish/subscribe attributes: for details see, Queue manager parameters

Table 135. Publish/subscribe configuration parameters

| Description                                               | MQSC parameter name |
|-----------------------------------------------------------|---------------------|
| Command message retry count                               | PSRTYCNT            |
| Discard undeliverable command input message               | PSNPMSG             |
| Behavior following undeliverable command response message | PSNPRES             |
| Process command messages under syncpoint                  | PSSYNCPT            |

The Broker stanza is used to manage the following configuration settings:

- `PersistentPublishRetry=yes | force`

If you specify `Yes`, then if a publication of a persistent message through the queued publish/subscribe interface fails, and no negative reply was requested, the publish operation is retried.

If you requested a negative response message, the negative response is sent and no further retry occurs.

If you specify `Force`, then if a publication of a persistent message through the queued publish/subscribe interface fails, the publish operation is retried until the it is successfully processed. No negative response is sent.

- `NonPersistentPublishRetry=yes | force`

If you specify `Yes`, then if a publication of a non-persistent message through the queued publish/subscribe interface fails, and no negative reply was requested, the publish operation is retried.

If you requested a negative response message, the negative response is sent and no further retry occurs.

If you specified `Force`, then if a publication of a non-persistent message through the queued publish/subscribe interface fails, the publish operation is retried until it is successfully processed. No negative response is sent.

**Note:** If you want to enable this functionality for non-persistent messages, then as well as setting the `NonPersistentPublishRetry` value you must also ensure that the queue manager attribute `PSSYNCPT` is set to `Yes`.

Doing this might also have an impact on the performance of processing non-persistent publications as the `MQGET` from the `STREAM` queue now occurs under syncpoint.

- `PublishBatchSize=number`

The broker normally processes publish messages within syncpoint. It can be inefficient to commit each publication individually, and in some circumstances the broker can process multiple publish messages in a single unit of work. This parameter specifies the maximum number of publish messages that can be processed in a single unit of work

The default value for `PublishBatchSize` is 5.

- `PublishBatchInterval=number`

The broker normally processes publish messages within syncpoint. It can be inefficient to commit each publication individually, and in some circumstances the broker can process multiple publish messages in a single unit of work. This parameter specifies the maximum time (in milliseconds) between the first message in a batch and any subsequent publication included in the same batch.

A batch interval of 0 indicates that up to `PublishBatchSize` messages can be processed, provided that the messages are available immediately.

The default value for `PublishBatchInterval` is zero.

## Procedure

Use IBM MQ Explorer, programmable commands, or the `runmqsc` command to alter the queue manager attributes that control the behavior of publish/subscribe.

## Example

```
ALTER QMGR PSNPRES(SAFE)
```

## Starting queued publish/subscribe

### Before you begin

Read the description of `PSMODE` to understand the three modes of publish/subscribe:

- COMPAT
- DISABLED
- ENABLED

### About this task

Set the `QMGR PSMODE` attribute to start either the queued publish/subscribe interface (also known as the broker), or the publish/subscribe engine (also known as Version 7 publish/subscribe) or both. To start queued publish/subscribe you need to set `PSMODE` to `ENABLED`. The default is `ENABLED`.

## Procedure

Use IBM MQ Explorer or the `runmqsc` command to enable the queued publish/subscribe interface if the interface is not already enabled.

## Example

```
ALTER QMGR PSMODE (ENABLED)
```

## What to do next

IBM MQ processes queued publish/subscribe commands and publish/subscribe Message Queue Interface (MQI) calls.

# Stopping queued publish/subscribe

## Before you begin

Read the description of PSMODE to understand the three modes of publish/subscribe:

- COMPAT
- DISABLED
- ENABLED

## About this task

Set the QMGR PSMODE attribute to stop either the queued publish/subscribe interface (also known as the broker), or the publish/subscribe engine (also known as Version 7 publish/subscribe) or both. To stop queued publish/subscribe you need to set PSMODE to COMPAT. To stop the publish/subscribe engine entirely, set PSMODE to DISABLED.

## Procedure

Use IBM MQ Explorer or the `runmqsc` command to disable the queued publish/subscribe interface.

## Example

```
ALTER QMGR PSMODE (COMPAT)
```

## Adding a stream

You can add streams manually to allow for data isolation between applications, or to allow inter-operation with Version 6 publish/subscribe hierarchies.

## Before you begin

Familiarize yourself with the way publish/subscribe streams operate. See Streams and topics.

## About this task

Use PCF command, `runmqsc`, or MQ Explorer to do these steps.

**Note:** You can perform steps 1 and 2 in any order. Only perform step 3 after steps 1 and 2 have both been completed.

## Procedure

1. Define a local queue with the same name as the Version 6 stream.
2. Define a local topic with the same name as the Version 6 stream.
3. Add the name of the queue to the namelist, `SYSTEM.QPUBSUB.QUEUE.NAMELIST`
4. Repeat for all queue managers at Version 7.1 or above that are in the publish/subscribe hierarchy.

## Adding 'Sport'

In the example of sharing the stream 'Sport', Version 6 and Version 7.1 queue managers are working in the same publish/subscribe hierarchy. The Version 6 queue managers share a stream called 'Sport'. The example shows how to create a queue and a topic on Version 7.1 queue managers called 'Sport', with a topic string 'Sport' that is shared with the Version 6 stream 'Sport'.

A Version 7.1 publish application, publishing to topic 'Sport', with topic string 'Soccer/Results', creates the resultant topic string 'Sport/Soccer/Results'. On Version 7.1 queue managers, subscribers to topic 'Sport', with topic string 'Soccer/Results' receive the publication.

On Version 6 queue managers, subscribers to stream 'Sport', with topic string 'Soccer/Results' receive the publication.

```
runmqsc QM1
5724-H72 (C) Copyright IBM Corp. 1994, 2008. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM1.
define qlocal('Sport')
 1 : define qlocal('Sport')
AMQ8006: IBM MQ queue created.
define topic('Sport') topicstr('Sport')
 2 : define topic('Sport') topicstr('Sport')
AMQ8690: IBM MQ topic created.
alter namelist(SYSTEM.QPUBSUB.QUEUE.NAMELIST) NAMES('Sport', 'SYSTEM.BROKER.DEFAULT.STREAM', 'SYSTEM.BROKER.ADMIN.STREAM')
 3 : alter namelist(SYSTEM.QPUBSUB.QUEUE.NAMELIST) NAMES('Sport', 'SYSTEM.BROKER.DEFAULT.STREAM', 'SYSTEM.BROKER.ADMIN.STREAM')
AMQ8551: IBM MQ namelist changed.
```

**Note:** You need both to provide the existing names in the namelist object, as well as the new names that you are adding, to the **alter namelist** command.

## What to do next

Information about the stream is passed to other brokers in the hierarchy.

If a broker is Version 6, administer it as a Version 6 broker. That is, you have a choice of creating the stream queue manually, or letting the broker create the stream queue dynamically when it is needed. The queue is based on the model queue definition, `SYSTEM.BROKER.MODEL.STREAM`.

If a broker is Version 7.1, you must configure each Version 7.1 queue manager in the hierarchy manually.

## Deleting a stream

You can delete a stream from an IBM MQ Version 7.1, or later, queue manager.

### Before you begin

Before deleting a stream you must ensure that there are no remaining subscriptions to the stream and quiesce all applications that use the stream. If publications continue to flow to a deleted stream, it takes a lot of administrative effort to restore the system to a cleanly working state.

### Procedure

1. Find all the connected brokers that host this stream.
2. Cancel all subscriptions to the stream on all the brokers.
3. Remove the queue (with the same name as the stream) from the namelist, `SYSTEM.QPUBSUB.QUEUE.NAMELIST`.
4. Delete or purge all the messages from the queue with the same name as the stream.
5. Delete the queue with the same name as the stream.
6. Delete the associated topic object.

## What to do next

Repeat steps 3 to 5 on all the other connected Version 7.1, or later, queue managers hosting the stream.

## Adding a subscription point

How to extend an existing queued publish/subscribe application that you have migrated from an earlier version of IBM Integration Bus with a new subscription point.

### Before you begin

1. Check that the subscription point is not already defined in `SYSTEM.QPUBSUB.SUBPOINT.NAMELIST`.
2. Check if there is a topic object or a topic string with the same name as the subscription point.

### About this task

IBM WebSphere MQ Version 7.1, or later, applications do not use subscription points, but they can interoperate with existing applications that do, using the subscription point migration mechanism.

**Important:** The subscription point migration mechanism has been removed from IBM MQ Version 8.0. If you need to migrate your existing applications, you must carry out the procedures described in the documentation for your version of the product, before you migrate to the latest version.

Subscription points do not work with queued publish/subscribe programs that use MQRFH1 headers, which have been migrated from IBM MQ Version 6, or earlier.

There is no need to add subscription points to use integrated publish/subscribe applications written for IBM MQ Version 7.1, or later.

### Procedure

1. Add the name of the subscription point to `SYSTEM.QPUBSUB.SUBPOINT.NAMELIST`.
  - On z/OS, the **NLTYPE** is `NONE`, the default.
  - Repeat the step on every queue manager that is connected in the same publish/subscribe topology.
2. Add a topic object, preferably giving it the name of the subscription point, with a topic string matching the name of the subscription point.
  - If the subscription point is in a cluster, add the topic object as a cluster topic on the cluster topic host.
  - If a topic object exists with the same topic string as the name of the subscription point, use the existing topic object. You must understand the consequences of the subscription point reusing an existing topic. If the existing topic is part of an existing application, you must resolve the collision between two identically named topics.
  - If a topic object exists with the same name as the subscription point, but a different topic string, create a topic with a different name.
3. Set the **Topic** attribute `WILDCARD` to the value `BLOCK`.

Blocking subscriptions to `#` or `*` isolates wildcard subscriptions to subscription points, see [Wildcards and subscription points](#).
4. Set any attributes that you require in the topic object.

### Example

The example shows a `runmqsc` command file that adds two subscription points, `USD` and `GBP`.

```
DEFINE TOPIC(USD) TOPICSTR(USD)
DEFINE TOPIC(GBP) TOPICSTR(GBP) WILDCARD(BLOCK)
ALTER NL(SYSTEM.QPUBSUB.SUBPOINT.NAMELIST) NAMES(SYSTEM.BROKER.DEFAULT.SUBPOINT, USD, GBP)
```

#### Note:

1. Include the default subscription point in the list of subscription points added using the **ALTER** command. **ALTER** deletes existing names in the `namelist`.



2. Define the topics before altering the namelist. The queue manager only checks the namelist when the queue manager starts and when the namelist is altered.

## Configuring distributed publish/subscribe networks

Queue managers that are connected together into a distributed publish/subscribe topology share a common federated topic space. Subscriptions created on one queue manager can receive messages published by an application connected to another queue manager in the topology.

You can control the extent of topic spaces created by connecting queue managers together in clusters or hierarchies. In a publish/subscribe cluster, a topic object must be 'clustered' for each branch of the topic space that is to span the cluster. In a hierarchy, each queue manager must be configured to identify its 'parent' in the hierarchy.

You can further control the flow of publications and subscriptions within the topology by choosing whether each publication and subscription is either local or global. Local publications and subscriptions are not propagated beyond the queue manager to which the publisher or subscriber is connected.

### Related information:

Distributed publish/subscribe networks

Publication scope

Subscription scope

Topic spaces

Defining cluster topics

## Configuring a publish/subscribe cluster

Define a topic on a queue manager. To make the topic a cluster topic, set the **CLUSTER** property. To choose the routing to use for publications and subscriptions for this topic, set the **CLROUTE** property.

### Before you begin

Some cluster configurations cannot accommodate the overheads of direct routed publish/subscribe. Before you use this configuration, explore the considerations and options detailed in Designing publish/subscribe clusters.

For changes to a cluster to be propagated throughout the cluster, at least one full repository must always be available. Ensure that your repositories are available before starting this task.

See also Routing for publish/subscribe clusters: Notes on behavior.

Scenario:

- The INVENTORY cluster has been set up as described in “Adding a queue manager to a cluster” on page 924. It contains three queue managers; LONDON and NEWYORK both hold full repositories, PARIS holds a partial repository.

### About this task

When you define a topic on a queue manager in a cluster, you need to specify whether the topic is a cluster topic, and (if so) the routing within the cluster for publications and subscriptions for this topic. To make the topic a cluster topic, you configure the **CLUSTER** property on the TOPIC object with the name of the cluster. By defining a cluster topic on a queue manager in the cluster, you make the topic available to the whole cluster. To choose the message routing to use within the cluster, you set the **CLROUTE** property on the TOPIC object to one of the following values:

- **DIRECT**
- **TOPICHOST**

By default, topic routing is **DIRECT**. This was the only option prior to IBM MQ Version 8.0. When you configure a direct routed clustered topic on a queue manager, all queue managers in the cluster become aware of all other queue managers in the cluster. When performing publish and subscribe operations, each queue manager can connect direct to any other queue manager in the cluster. See Direct routed publish/subscribe clusters.

From IBM MQ Version 8.0, you can instead configure topic routing as **TOPICHOST**. When you use topic host routing, all queue managers in the cluster become aware of the cluster queue managers that host the routed topic definition (that is, the queue managers on which you have defined the topic object). When performing publish and subscribe operations, queue managers in the cluster connect only to these topic host queue managers, and not directly to each other. The topic host queue managers are responsible for routing publications from queue managers on which publications are published to queue managers with matching subscriptions. See Topic host routed publish/subscribe clusters.

**Note:** After a topic object has been clustered (through setting the **CLUSTER** property) you cannot change the value of the **CLROUTE** property. The object must be un-clustered (**CLUSTER** set to ' ') before you can change the value. Un-clustering a topic converts the topic definition to a local topic, which results in a period during which publications are not delivered to subscriptions on remote queue managers; this should be considered when performing this change. See The effect of defining a non-cluster topic with the same name as a cluster topic from another queue manager. If you try to change the value of the **CLROUTE** property while it is clustered, the system generates an MQRCCF\_CLROUTE\_NOT\_ALTERABLE exception.

## Procedure

1. Choose a queue manager to host your topic.

Any cluster queue manager can host a topic. Choose one of the three queue managers ( LONDON, NEWYORK or PARIS) and configure the properties of the TOPIC object. If you plan to use direct routing, it makes no operational difference which queue manager you choose. If you plan to use topic host routing, the chosen queue manager has additional responsibilities for routing publications. Therefore, for topic host routing, choose a queue manager that is hosted on one of your more powerful systems and has good network connectivity.

2. Define a topic on a queue manager.

To make the topic a cluster topic, include the cluster name when you define the topic, and set the routing that you want to use for publications and subscriptions for this topic. For example, to create a direct routing cluster topic on the LONDON queue manager, create the topic as follows:

```
DEFINE TOPIC(INVENTORY) TOPICSTR('/INVENTORY') CLUSTER(INVENTORY) CLROUTE(DIRECT)
```

By defining a cluster topic on a queue manager in the cluster, you make the topic available to the whole cluster.

For more information about using **CLROUTE**, see DEFINE TOPIC (CLROUTE) and Routing for publish/subscribe clusters: Notes on behavior.

## Results

The cluster is ready to receive publications and subscriptions for the topic.

## What to do next

If you have configured a topic host routed publish/subscribe cluster, you will probably want to add a second topic host for this topic. See “Adding extra topic hosts to a topic host routed cluster” on page 1012.

If you have several separate publish/subscribe clusters, for example because your organization is geographically dispersed, you might want to propagate some cluster topics into all the clusters. You can do this by connecting the clusters in a hierarchy. See “Combining the topic spaces of multiple clusters” on page 1016. You can also control which publications flow from one cluster to another. See “Combining and

isolating topic spaces in multiple clusters” on page 1018.

**Related information:**

Designing publish/subscribe clusters

Distributed publish/subscribe troubleshooting

Inhibiting clustered publish/subscribe

## Moving a cluster topic definition to a different queue manager

For either topic host routed or direct routed clusters, you might need to move a cluster topic definition when decommissioning a queue manager, or because a cluster queue manager has failed or is unavailable for a significant period of time.

### About this task

You can have multiple definitions of the same cluster topic object in a cluster. This is a normal state for a topic host routed cluster, and an unusual state for a direct routed cluster. For more information, see Multiple cluster topic definitions of the same name.

To move a cluster topic definition to a different queue manager in the cluster without interrupting the flow of publications, complete the following steps. The procedure moves a definition from queue manager QM1 to queue manager QM2.

### Procedure

1. Create a duplicate of the cluster topic definition on QM2.

For direct routing, set all the attributes to match the definition of QM1.

For topic host routing, initially define the new topic host as PUB(DISABLED). This allows QM2 to learn of the subscriptions in the cluster, but not to start routing publications.

2. Wait for information to be propagated through the cluster.

Wait for the new cluster topic definition to be propagated by the full repository queue managers to all queue managers in the cluster. Use the **DISPLAY CLUSTER** command to display the cluster topics on each cluster member, and check for a definition originating from QM2.

For topic host routing, wait for the new topic host on QM2 to learn of all subscriptions. Compare the proxy subscriptions known to QM2 and those known to QM1. One way to view the proxy subscriptions on a queue manager is to issue the following **runmqsc** command:

```
DISPLAY SUB(*) SUBTYPE(PROXY)
```

3. For topic host routing, redefine the topic host on QM2 as PUB(ENABLED), then redefine the topic host on QM1 as PUB(DISABLED).

Now that the new topic host on QM2 has learnt of all subscriptions on other queue managers, the topic host can start routing publications.

By using the PUB(DISABLED) setting to quiesce message traffic through QM1, you ensure that no publications are in train through QM1 when you delete the cluster topic definition.

4. Delete the cluster topic definition from QM1.

You can only delete the definition from QM1 if the queue manager is available. Otherwise, you must run with both definitions in existence until QM1 is restarted or forcibly removed.

If QM1 remains unavailable for a long time, and during that time you need to modify the clustered topic definition on QM2, the QM2 definition is newer than the QM1 definition, and therefore usually prevails.

During this period, if there are differences between the definitions on QM1 and QM2, errors are written to the error logs of both queue managers, alerting you to the conflicting cluster topic definition.

If QM1 is never going to return to the cluster, for example because of unexpected decommissioning following a hardware failure, as a last resort you can use the `RESET CLUSTER` command to forcibly eject the queue manager. **RESET CLUSTER** automatically deletes all topic objects hosted on the target queue manager.

## Adding extra topic hosts to a topic host routed cluster

In a topic host routed publish/subscribe cluster, multiple queue managers can be used to route publications to subscriptions by defining the same clustered topic object on those queue managers. This can be used to improve availability and workload balancing. When you add an extra topic host for the same cluster topic object, you can use the **PUB** parameter to control when publications begin to be routed through the new topic host.

### Before you begin

Defining the same cluster topic object on several queue managers is only functionally useful for a topic host routed cluster. Defining multiple matching topics in a direct routed cluster does not change its behavior. This task only applies to topic host routed clusters.

This task assumes that you have read the article *Multiple cluster topic definitions of the same name*, especially the following sections:

- Multiple cluster topic definitions in a topic host routed cluster
- Special handling for the **PUB** parameter

### About this task

When a queue manager is made a routed topic host, it must first learn of the existence of all related topics that have been subscribed to in the cluster. If publications are being published to those topics at the time that an additional topic host is added, and a publication is routed to the new host before that host has learned of the existence of subscriptions on other queue managers in the cluster, then the new host does not forward that publication to those subscriptions. This causes subscriptions to miss publications.

Publications are not routed through topic host queue managers that have explicitly set the cluster topic object **PUB** parameter to `DISABLED`, so you can use this setting to ensure that no subscriptions miss publications during the process of adding an extra topic host.

**Note:** While a queue manager hosts a cluster topic that has been defined as `PUB(DISABLED)`, publishers connected to that queue manager cannot publish messages, and matching subscriptions on that queue manager do not receive publications published on other queue managers in the cluster. For this reason, careful consideration must be given to defining topic host routed topics on queue managers where subscriptions exist and publishing applications connect.

### Procedure

1. Configure a new topic host, and initially define the new topic host as `PUB(DISABLED)`.  
This allows the new topic host to learn of the subscriptions in the cluster, but not to start routing publications.  
For information about configuring a topic host, see “Configuring a publish/subscribe cluster” on page 1009.
2. Determine when the new topic host has learned of all subscriptions.  
To do this, compare the proxy subscriptions known to the new topic host and those known to the existing topic host. One way to view the proxy subscriptions is to issue the following **runmqsc** command: `DISPLAY SUB(*) SUBTYPE(PROXY)`
3. Redefine the new topic host as `PUB(ENABLED)`.

After the new topic host has learned of all subscriptions on other queue managers, the topic can start routing publications.

## Combining publication and subscription scopes

In IBM MQ versions 7 onwards, publication and subscription scope work independently to determine the flow of publications between queue managers.

Publications can flow to all queue managers that are connected in a publish/subscribe topology, or only to the local queue manager. Similarly for proxy subscriptions. Which publications match a subscription is governed by the combination of these two flows.

Publications and subscriptions can both be scoped to QMGR or ALL. If a publisher and a subscriber are both connected to the same queue manager, scope settings do not affect which publications the subscriber receives from that publisher.

If the publisher and subscriber are connected to different queue managers, both settings must be ALL to receive remote publications.

Suppose publishers are connected to different queue managers. If you want a subscriber to receive publications from any publisher, set the subscription scope to ALL. You can then decide, for each publisher, whether to limit the scope of its publications to subscribers local to the publisher.

Suppose subscribers are connected to different queue managers. If you want the publications from a publisher to be sent to all the subscribers, set the publication scope to ALL. If you want a subscriber to receive publications only from a publisher connected to the same queue manager, set the subscription scope to QMGR.

### Example: football results service

Suppose you are a member team in a football league. Each team has a queue manager connected to all the other teams in a publish/subscribe cluster.

The teams publish the results of all the games played on their home ground using the topic, *Football/result/Home team name/Away team name*. The strings in italics are variable topic names, and the publication is the result of the match.

Each club also republishes the results just for the club using the topic string *Football/myteam/Home team name/Away team name*.

Both topics are published to the whole cluster.

The following subscriptions have been set up by the league so that fans of any team can subscribe to the results in three interesting ways.

Notice that you can set up cluster topics with SUBSCOPE(QMGR). The topic definitions are propagated to each member of the cluster, but the scope of the subscription is just the local queue manager. Thus subscribers at each queue manager receive different publications from the same subscription.

#### Receive all results

```
DEFINE TOPIC(A) TOPICSTR('Football/result/') CLUSTER SUBSCOPE(ALL)
```

#### Receive all home results

```
DEFINE TOPIC(B) TOPICSTR('Football/result/') CLUSTER SUBSCOPE(QMGR)
```

Because the subscription has QMGR scope, only results published at the home ground are matched.

#### Receive all my teams results

```
DEFINE TOPIC(C) TOPICSTR('Football/myteam/') CLUSTER SUBSCOPE(QMGR)
```

Because the subscription has QMGR scope, only the local team results, which are republished locally, are matched.

**Related information:**

Distributed publish/subscribe networks  
Publication scope  
Subscription scope

**Combining topic spaces in publish/subscribe networks**

Combine the topic space of a queue manager with other queue managers in a publish/subscribe cluster or hierarchy. Combine publish/subscribe clusters, and publish/subscribe clusters with hierarchies.

You can create different publish/subscribe topic spaces by using the building blocks of **CLUSTER**, **PUBSCOPE** and **SUBSCOPE** attributes, publish/subscribe clusters, and publish/subscribe hierarchies.

Starting from the example of scaling up from a single queue manager to a publish/subscribe cluster, the following scenarios illustrate different publish/subscribe topologies.

**Related information:**

Distributed publish/subscribe networks  
Topic spaces  
Defining cluster topics

**Creating a single topic space in a publish/subscribe cluster:**

Scale up a publish/subscribe system to run on multiple queue managers. Use a publish/subscribe cluster to provide each publisher and subscriber with a single identical topic space.

**Before you begin**

You have implemented a publish/subscribe system on a single version 7 queue manager.

Always create topic spaces with their own root topics, rather than relying on inheriting the attributes of `SYSTEM.BASE.TOPIC`. If you scale your publish/subscribe system up to a cluster, you can define your root topics as cluster topics, on the cluster topic host, and then all your topics are shared throughout the cluster.

**About this task**

You now want to scale the system up to support more publishers and subscribers and have every topic visible throughout the cluster.

**Procedure**

1. Create a cluster to use with the publish/subscribe system. If you have an existing traditional cluster, for performance reasons it is better to set up a new cluster for the new publish subscribe system. You can use the same servers for the cluster repositories of both clusters
2. Choose one queue manager, possibly one of the repositories, to be the cluster topic host.
3. Ensure every topic that is to be visible throughout the publish/subscribe cluster resolves to an administrative topic object. Set the **CLUSTER** attribute naming the publish/subscribe cluster.

**What to do next**

Connect publisher and subscriber applications to any queue managers in the cluster.

Create administrative topic objects that have the **CLUSTER** attribute. The topics are also propagated throughout the cluster. Publisher and subscriber programs use the administrative topics so that their behavior is not altered by being connected to different queue managers in the cluster

If you need `SYSTEM.BASE.TOPIC` to act like a cluster topic on every queue manager, you need to modify it on every queue manager.

**Related information:**

Distributed publish/subscribe networks

Topic spaces

Defining cluster topics

**Adding a version 7 or later queue manager to existing Version 6 topic spaces:**

Extend an existing Version 6 publish/subscribe system to interoperate with a version 7 or later queue manager, sharing the same topic spaces.

**Before you begin**

You have an existing Version 6 publish/subscribe system.

You have installed IBM MQ version 7 or later on a new server and configured a queue manager.

**About this task**

You want to extend your existing Version 6 publish/subscribe system to work with version 7 or later queue managers.

You have decided to stabilize development of the Version 6 publish/subscribe system that uses the queued publish/subscribe interface. You intend to add extensions to the system using the version 7 or later MQI. You have no plans now to rewrite the queued publish/subscribe applications.

You intend to upgrade the Version 6 queue managers to version 7 or later in the future. For now, you are continuing to run the existing queued publish/subscribe applications on the version 7 or later queue managers.

**Procedure**

1. Create one set of sender-receiver channels to connect the version 7 or later queue manager with one of the Version 6 queue managers in both directions.
2. Create two transmission queues with the names of the target queue managers. Use queue manager aliases if you cannot use the name of the target queue manager as the transmission queue name for some reason.
3. Configure the transmission queues to trigger the sender channels.
4. If the Version 6 publish/subscribe system uses streams, add the streams to the version 7 or later queue manager as described in Adding a stream.
5. Check the version 7 or later queue manager **PSMODE** is set to **ENABLE**.
6. Alter its **PARENT** attribute to refer to one of the Version 6 queue managers.
7. Check the status of the parent-child relationship between the queue managers is active in both directions.

**What to do next**

Once you have completed the task, both the Version 6 and version 7 or later queue manager share the same topic spaces. For example, you can do all the following tasks.

- Exchange publications and subscriptions between Version 6 and version 7 or later queue managers.
- Run your existing Version 6 publish/subscribe programs on the version 7 or later queue manager.
- View and modify the topic space on either the Version 6 or version 7 or later queue manager.
- Write version 7 or later publish/subscribe applications and run them on the version 7 or later queue manager.
- Create new publications and subscriptions with the version 7 or later applications and exchange them with Version 6 applications.

**Related information:**

Distributed publish/subscribe networks

Topic spaces

Defining cluster topics

**Combining the topic spaces of multiple clusters:**

Create topic spaces that span multiple clusters. Publish to a topic in one cluster and subscribe to it in another.

**Before you begin**

This task assumes that you have existing direct routed publish/subscribe clusters, and you want to propagate some cluster topics into all the clusters.

**Note:** You cannot do this for topic host routed publish/subscribe clusters.

**About this task**

To propagate publications from one cluster to another, you need to join the clusters together in a hierarchy; see Figure 146 on page 1017. The hierarchical connections propagate subscriptions and publications between the connected queue managers, and the clusters propagate cluster topics within each cluster, but not between clusters.

The combination of these two mechanisms propagates cluster topics between all the clusters. You need to repeat the cluster topic definitions in each cluster.



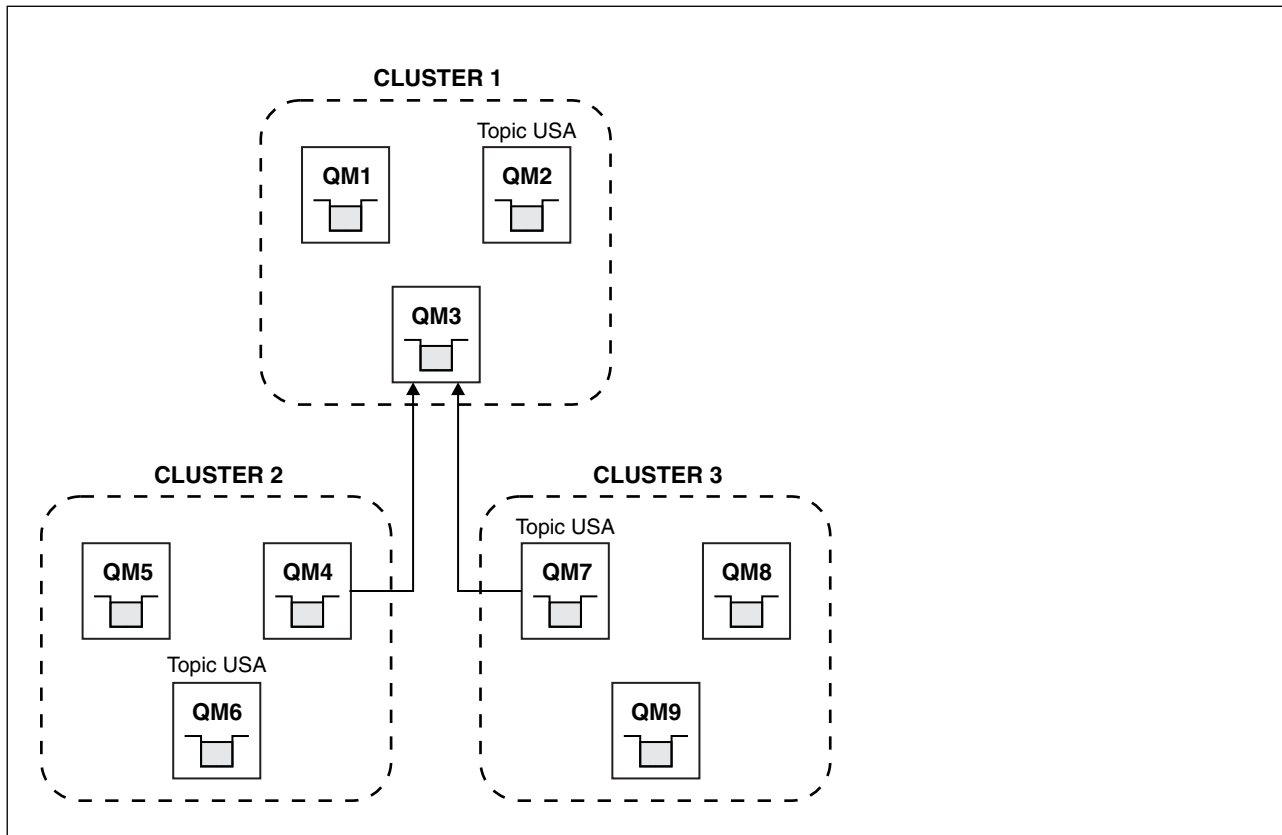


Figure 146. Connecting clusters using hierarchies

The following steps connect the clusters into a hierarchy.

#### Procedure

1. Create two sets of sender-receiver channels to connect QM3 and QM4, and QM3 and QM7, in both directions. You must use traditional sender-receiver channels and transmission queues, rather than a cluster, to connect a hierarchy.
2. Create three transmission queues with the names of the target queue managers. Use queue manager aliases if you cannot use the name of the target queue manager as the transmission queue name for some reason.
3. Configure the transmission queues to trigger the sender channels.
4. Check the **PSMODE** of QM3, QM4 and QM7 is set to ENABLE.
5. Alter the **PARENT** attribute of QM4 and QM7 to QM3.
6. Check the status of the parent-child relationship between the queue managers is active in both directions.
7. Create the administrative topic USA with the attribute **CLUSTER** ( 'CLUSTER 1' ), **CLUSTER** ( 'CLUSTER 2' ), and **CLUSTER** ( 'CLUSTER 3' ) on each of the three cluster topic host queue managers in clusters 1, 2 and 3. The cluster topic host does not need to be a hierarchically connected queue manager.

#### What to do next

You can now publish or subscribe to the cluster topic USA in Figure 146. The publications subscriptions flow to publishers and subscribers in all three clusters.

Suppose that you did not create USA as a cluster topic in the other clusters. If USA is only defined on QM7, then publications and subscriptions to USA are exchanged between QM7, QM8, QM9, and QM3. Publishers and

subscribers running on QM7, QM8, QM9 inherit the attributes of the administrative topic USA. Publishers and subscribers on QM3 inherit the attributes of SYSTEM.BASE.TOPIC on QM3.

See also “Combining and isolating topic spaces in multiple clusters.”

**Related information:**

Distributed publish/subscribe networks

Topic spaces

Defining cluster topics

**Combining and isolating topic spaces in multiple clusters:**

Isolate some topic spaces to a specific cluster, and combine other topic spaces to make them accessible in all the connected clusters.

**Before you begin**

Examine the topic “Combining the topic spaces of multiple clusters” on page 1016. It might be sufficient for your needs, without adding an additional queue manager as a bridge.

**Note:** You can only complete this task using direct routed publish/subscribe clusters. You cannot do this using topic host routed clusters.

**About this task**

A potential improvement on the topology shown in Figure 146 on page 1017 in “Combining the topic spaces of multiple clusters” on page 1016 is to isolate cluster topics that are not shared across all the clusters. Isolate clusters by creating a bridging queue manager that is not in any of the clusters; see Figure 147 on page 1019. Use the bridging queue manager to filter which publications and subscriptions can flow from one cluster to another.

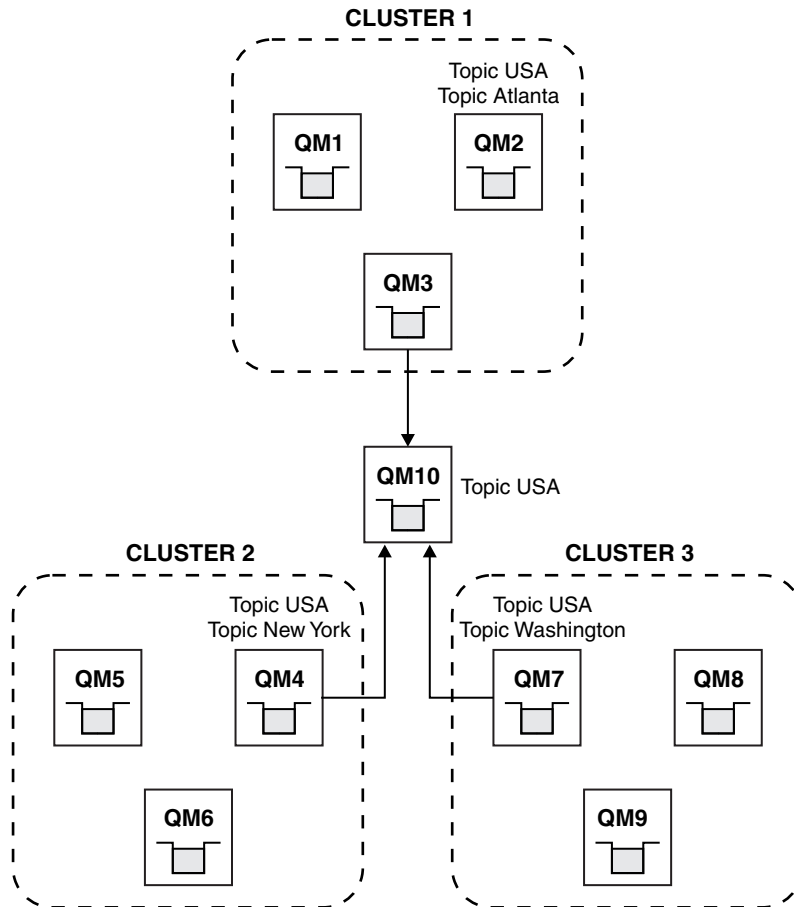


Figure 147. Bridged clusters

Use the bridge to isolate cluster topics that you do not want exposed across the bridge on the other clusters. In Figure 147, USA is a cluster topic shared in all the clusters, and Atlanta, New York and Washington are cluster topics that are shared only in one cluster each.

Model your configuration using the following procedure:

#### Procedure

1. Modify all the `SYSTEM.BASE.TOPIC` topic objects to have **SUBSCOPE** ( QMGR ) and **PUBSCOPE** ( QMGR ) on all the queue managers. No topics (even cluster topics) are propagated onto other queue managers unless you explicitly set **SUBSCOPE** ( ALL ) and **PUBSCOPE** ( ALL ) on the root topic of your cluster topics.
2. Define the topics on the three cluster topic host queue managers that you want to be shared in each cluster with the attributes **CLUSTER** ( *clustername* ), **SUBSCOPE** ( ALL ) and **PUBSCOPE** ( ALL ). If you want some cluster topics shared between all the clusters, define the same topic in each of the clusters. Use the cluster name of each cluster as the cluster attribute.
3. For the cluster topics you want shared between all the clusters, define the topics again on the bridge queue manager ( QM10 ), with the attributes **SUBSCOPE** ( ALL ), and **PUBSCOPE** ( ALL ).

#### Example

In the example in Figure 147, only topics that inherit from USA propagate between all three clusters.

## What to do next

Subscriptions for topics defined on the bridge queue manager with **SUBSCOPE** ( ALL ) and **PUBSCOPE** ( ALL ) are propagated between the clusters.

Subscriptions for topics defined within each cluster with attributes **CLUSTER** (*clustername*), **SUBSCOPE** ( ALL ) and **PUBSCOPE** ( ALL ) are propagated within each cluster.

Any other subscriptions are local to a queue manager.

### Related information:

Distributed publish/subscribe networks

Topic spaces

Defining cluster topics

Publication scope

Subscription scope

### Publishing and subscribing to topic spaces in multiple clusters:

Publish and subscribe to topics in multiple clusters using overlapped clusters. You can use this technique as long as the topic spaces in the clusters do not overlap.

### Before you begin

Create multiple traditional clusters with some queue managers in the intersections between the clusters.

### About this task

You might have chosen to overlap clusters for various different reasons.

1. You have a limited number of high availability servers, or queue managers. You decide to deploy all the cluster repositories, and cluster topic hosts to them.
2. You have existing traditional queue manager clusters that are connected using gateway queue managers. You want to deploy publish/subscribe applications to the same cluster topology.
3. You have a several self contained publish/subscribe applications. For performance reasons, it is better to keep publish/subscribe clusters small and separate from traditional clusters. You have decided to deploy the applications to different clusters. However, you also want to monitor all the publish/subscribe applications on one queue manager, as you have licensed only one copy of the monitoring application. This queue manager must have access to the publications to cluster topics in all the clusters.

By ensuring that your topics are defined in non-overlapping topic spaces, you can deploy the topics into overlapping publish/subscribe clusters, see Figure 148 on page 1021. If the topic spaces overlap, then deploying to overlapping clusters leads to problems.

Because the publish/subscribe clusters overlap you can publish and subscribe to any of the topic spaces using the queue managers in the overlap.

Figure 148. Overlapping clusters, non-overlapping topic spaces

### Procedure

Create a means of ensuring that topic spaces do not overlap. For example, define a unique root topic for each of the topic spaces. Make the root topics cluster topics.

1. DEFINE TOPIC(B) TOPICSTR('B') CLUSTER('CLUSTER 1') ...
2. DEFINE TOPIC(C) TOPICSTR('C') CLUSTER('CLUSTER 2') ...

### Example

In Figure 148 publishers and subscriber connected to QM3 can publish or subscribe to  $T_B$  or  $T_C$

### What to do next

Connect publishers and subscribers that use topics in both clusters to queue managers in the overlap.

Connect publishers and subscribers that must only use topics in a specific cluster to queue managers not in the overlap.

### Related information:

Distributed publish/subscribe networks

Topic spaces

Defining cluster topics

## Connecting a queue manager to a publish/subscribe hierarchy

You connect the child queue manager to the parent queue manager in the hierarchy. If the child queue manager is already a member of another hierarchy or cluster, then this connection joins the hierarchies together, or joins the cluster to the hierarchy.

### Before you begin

1. Queue managers in a publish/subscribe hierarchy must have unique queue manager names.
2. A publish/subscribe hierarchy relies on the “queued publish/subscribe” queue manager feature. This must be enabled on both the parent and the child queue managers. See Starting queued publish/subscribe.
3. The publish/subscribe relationship relies on queue manager sender and receiver channels. There are two ways to establish the channels:
  - Add both the parent and child queue managers to a IBM MQ cluster. See Adding a queue manager to a cluster.
  - Establish a sender/receiver channel pair from the child queue manager to the parent and from the parent to the child. Each channel either needs to use a transmission queue with the same name as the target queue manager, or a queue manager alias with the same name as the target queue manager. For more information about how to establish a point-to-point channel connection, see IBM MQ distributed queuing techniques.

For examples that configure a hierarchy over each type of channel configuration, see the following set of publish/subscribe hierarchy scenarios:

- Scenario 1: Using point-to-point channels with queue manager name alias
- Scenario 2: Using point-to-point channels with same name for transmission queue and remote queue manager
- Scenario 3: Using a cluster channel to add a queue manager

## About this task

Use the ALTER QMGR PARENT (*PARENT\_NAME*) **runmqsc** command to connect children to parents. This configuration is performed on the child queue manager, where *PARENT\_NAME* is the name of the parent queue manager.

## Procedure

```
ALTER QMGR PARENT(PARENT_NAME)
```

## Example

The first example shows how to attach queue manager QM2 as a child of QM1, then query QM2 to confirm it has successfully become a child with a **STATUS** of ACTIVE:

```
C:>runmqsc QM2
5724-H72 (C) Copyright IBM Corp. 1994, 2008. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM2
alter qmgr parent(QM1)
 1 : alter qmgr parent(QM1)
AMQ8005: IBM MQ queue manager changed.
display pubsub all
 2 : display pubsub all
AMQ8723: Display pub/sub status details.
 QMNAME(QM2) TYPE(LOCAL)
 STATUS(ACTIVE)
AMQ8723: Display pub/sub status details.
 QMNAME(QM1) TYPE(PARENT)
 STATUS(ACTIVE)
```

The next example shows the result of querying QM1 for its connections:

```
C:\Documents and Settings\Admin>runmqsc QM1
5724-H72 (C) Copyright IBM Corp. 1994, 2008. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM1.
display pubsub all
 2 : display pubsub all
AMQ8723: Display pub/sub status details.
 QMNAME(QM1) TYPE(LOCAL)
 STATUS(ACTIVE)
AMQ8723: Display pub/sub status details.
 QMNAME(QM2) TYPE(CHILD)
 STATUS(ACTIVE)
```

If **STATUS** does not show as ACTIVE, check that the channels between the child and the parent are correctly configured and running. Check both queue manager error logs for possible errors.

## What to do next

By default, topics used by publishers and subscribers on one queue manager are shared with publishers and subscribers on the other queue managers in the hierarchy. Administered topics can be configured to control the level of sharing through use of the **SUBSCOPE** and **PUBSCOPE** topic properties. See Configuring distributed publish/subscribe networks.

**Related information:**

Streams and topics

DISPLAY PUBSUB

Publish/subscribe messaging

**Disconnecting a queue manager from a publish/subscribe hierarchy**

Disconnect a child queue manager from a parent queue manager in a publish/subscribe hierarchy.

**About this task**

Use the **ALTER QMGR** command to disconnect a queue manager from a broker hierarchy. You can disconnect a queue manager in any order at any time.

The corresponding request to update the parent is sent when the connection between the queue managers is running.

**Procedure**

```
ALTER QMGR PARENT('')
```

**Example**

```
C:\Documents and Settings\Admin>runmqsc QM2
5724-H72 (C) Copyright IBM Corp. 1994, 2008. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM2.
 1 : alter qmgr parent('')
AMQ8005: IBM MQ queue manager changed.
 2 : display pubsub type(child)
AMQ8147: IBM MQ object not found.
display pubsub type(parent)
 3 : display pubsub type(parent)
AMQ8147: IBM MQ object not found.
```

**What to do next**

You can delete any streams, queues and manually defined channels that are no longer needed.

---

**Configuring multiple installations**

When using multiple installations on the same system, you must configure the installations and queue managers.

This information applies to UNIX, Linux, and Windows.

Use the information in the following links to configure your installations:

- “Changing the primary installation” on page 1033
- “Associating a queue manager with an installation” on page 1034
- “Connecting applications in a multiple installation environment” on page 1024

**Related information:**

Choosing a primary installation

Multiple installations

Choosing an installation name

## Connecting applications in a multiple installation environment

On UNIX, Linux, and Windows systems, if IBM WebSphere MQ Version 7.1, or later, libraries are loaded, IBM MQ automatically uses the appropriate libraries without you needing to take any further action. IBM MQ uses libraries from the installation associated with the queue manager that the application connects to.

The following concepts are used to explain the way applications connect to IBM MQ:

### Linking

When the application is compiled, the application is linked to the IBM MQ libraries to get the function exports that are then loaded when the application runs.

### Loading

When the application is run, the IBM MQ libraries are located and loaded. The specific mechanism used to locate the libraries varies by operating system, and by how the application is built. For more information about how to locate and load libraries in a multiple installation environment, see “Loading IBM WebSphere MQ Version 7.1, or later version, libraries” on page 1026.

### Connecting

When the application connects to a running queue manager, for example, using a MQCONN or MQCONNX call, it connects using the loaded IBM MQ libraries.

When a server application connects to a queue manager, the loaded libraries must come from the installation associated with the queue manager. With multiple installations on a system, this restriction introduces new challenges when choosing the mechanism that the operating system uses to locate the IBM MQ libraries to load:

- When the **setmqm** command is used to change the installation associated with a queue manager, the libraries that need to be loaded change.
- When an application connects to multiple queue managers that are owned by different installations, multiple sets of libraries need to be loaded.

However, if IBM WebSphere MQ Version 7.1, or later, libraries, are located and loaded, IBM MQ then loads and uses the appropriate libraries without you needing to take any further action. When the application connects to a queue manager, IBM MQ loads libraries from the installation that the queue manager is associated with.



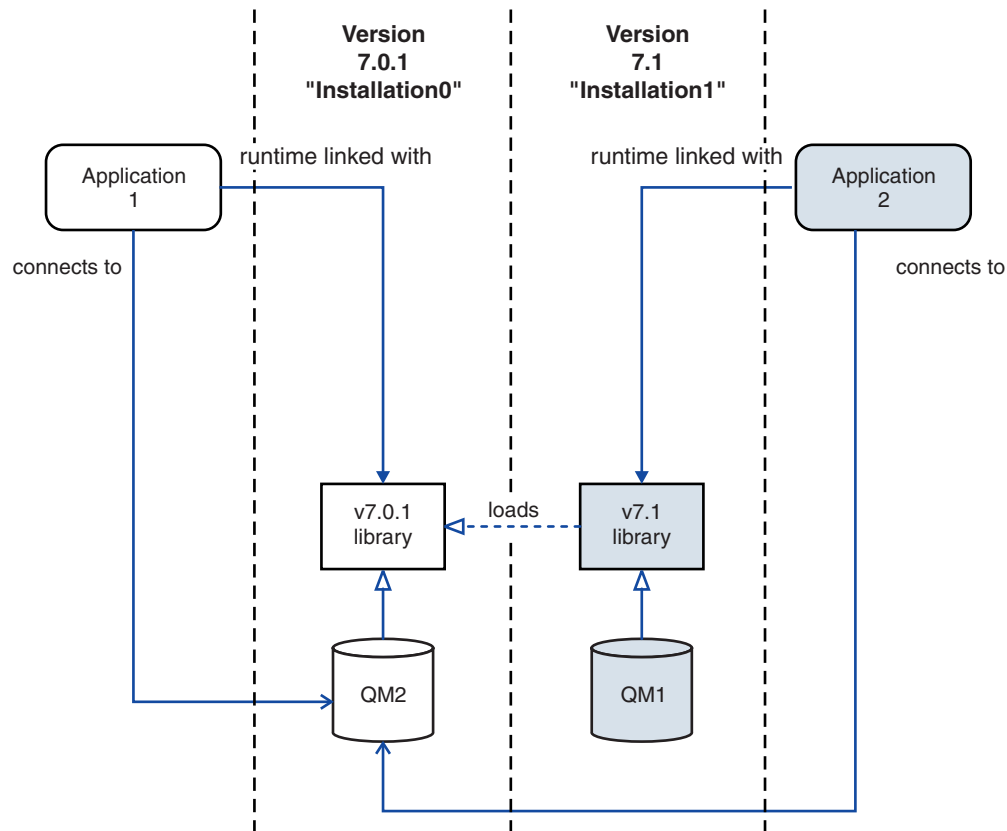


Figure 149. Connecting applications in a multiple installation environment

For example, Figure 149 shows a multiple installation environment with a Version 7.0.1 installation ( Installation0), and a Version 7.1 installation ( Installation1). Two applications are connected to these installations, but they load different library versions.

Application 1 directly loads a Version 7.0.1 library. When application 1 connects to QM2, the Version 7.0.1 libraries are used . If application 1 attempts to connect to QM1, or if QM2 is associated with Installation1, application 1 fails with a 2059 (080B) (RC2059): MQRC\_Q\_MGR\_NOT\_AVAILABLE error. The application fails because the Version 7.0.1 library is not capable of loading other library versions. That is, if Version 7.0.1 libraries are directly loaded, you cannot use a queue manager associated with an installation at a later version of IBM MQ.

Application 2 directly loads a Version 7.1 library. When application 2 connects to QM2, the Version 7.1 library then loads and uses the Version 7.0.1 library. If application 2 connects to QM1, or if QM2 is associated with Installation1, the Version 7.1 library is loaded, and the application works as expected.

Migration scenarios and connecting applications with multiple installations is considered in more detail in Multi-installation queue manager coexistence on UNIX, Linux, and Windows .

For more information about how to load IBM WebSphere MQ Version 7.1 libraries, see “Loading IBM WebSphere MQ Version 7.1, or later version, libraries” on page 1026.

## Support and restrictions

If any of the following Version 7.1, or later, libraries, are located and loaded, IBM MQ can automatically load and use the appropriate libraries:

- The C server libraries
- The C++ server libraries
- The XA server libraries
- The COBOL server libraries
- The COM+ server libraries
- .NET in unmanaged mode

IBM MQ also automatically loads and uses the appropriate libraries for Java and JMS applications in bindings mode.

There are a number of restrictions for applications using multiple installations. For more information, see “Restrictions for applications using multiple installations” on page 1029.

**Related concepts:**

“Associating a queue manager with an installation” on page 1034

When you create a queue manager, it is automatically associated with the installation that issued the **crtmqm** command. On UNIX, Linux, and Windows, you can change the installation associated with a queue manager using the **setmqm** command.

“Restrictions for applications using multiple installations” on page 1029

There are restrictions when using CICS server libraries, fast path connections, message handles, and exits in a multiple installation environment.

“Loading IBM WebSphere MQ Version 7.1, or later version, libraries”

When deciding how to load IBM MQ libraries, you need to consider a number of factors, including: your environment, whether you can change your existing applications, whether you want a primary installation, where IBM MQ is installed, and whether the location of IBM MQ is likely to change.

**Related tasks:**

“Changing the primary installation” on page 1033

You can use the **setmqinst** command to set or unset an installation as the primary installation.

**Related information:**

Choosing a primary installation

**Loading IBM WebSphere MQ Version 7.1, or later version, libraries**

When deciding how to load IBM MQ libraries, you need to consider a number of factors, including: your environment, whether you can change your existing applications, whether you want a primary installation, where IBM MQ is installed, and whether the location of IBM MQ is likely to change.

How IBM WebSphere MQ Version 7.1, or later version, libraries are located and loaded depends on your installation environment:

- On UNIX and Linux systems, if a copy of IBM WebSphere MQ Version 7.1, or later version, is installed in the default location, existing applications continue to work in the same way as previous versions. However, if the applications need symbolic links in `/usr/lib`, you must either select a Version 7.1, or later version, installation to be the primary installation, or manually create the symbolic links.
- If IBM WebSphere MQ Version 7.1, or later version, is installed in a non-default location, which is the case if IBM WebSphere MQ Version 7.0.1 is also installed, you might need to change your existing applications so that the correct libraries are loaded.

How IBM WebSphere MQ Version 7.1, or later version, libraries can be located and loaded also depends on how any existing applications are set up to load libraries. For more information about how libraries can be loaded, see “Operating system library loading mechanisms” on page 1028.

Optimally, you should ensure the IBM MQ library, that is loaded by the operating system, is the one with which the queue manager is associated.

The methods for loading IBM MQ libraries vary by platform, and each method has benefits and drawbacks.

Table 136. Benefits and drawbacks of the options for loading libraries

| Platform               | Option                                                                                                                                                                                                                                                         | Benefits                                                                                                                                                                                                                            | Drawbacks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNIX and Linux systems | <p>Set or change the embedded runtime search path (RPath) of the application.</p> <p>This option requires you to recompile and link the application. For more information about compiling and linking applications, see Building a procedural application.</p> | <ul style="list-style-type: none"> <li>• Scope of the change is clear.</li> </ul>                                                                                                                                                   | <ul style="list-style-type: none"> <li>• You must be able to recompile and link the application.</li> <li>• If the location of IBM MQ changes, you must change the RPath.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| UNIX and Linux systems | <p>Set the <code>LD_LIBRARY_PATH</code> environment variable (<code>LIBPATH</code> on AIX ), using <code>setmqenv</code>, or <code>crtmqenv</code>, with the <code>-k</code> or <code>-l</code> option.</p>                                                    | <ul style="list-style-type: none"> <li>• No changes to existing applications required.</li> <li>• Overrides embedded RPaths in an application.</li> <li>• Easy to change the variable if the location of IBM MQ changes.</li> </ul> | <ul style="list-style-type: none"> <li>• <code>setuid</code> and <code>setgid</code> applications, or applications built in other ways, might ignore <code>LD_LIBRARY_PATH</code> for security reasons.</li> <li>• Environment specific, so must be set in each environment where the application is run.</li> <li>• Possible impact on other applications that rely on <code>LD_LIBRARY_PATH</code>.</li> <li>• HP-UX: Options used when the application was compiled might disable the use of <code>LD_LIBRARY_PATH</code>. For more information, see Runtime linking considerations for HP-UX.</li> <li>• Linux: The compiler used to build the application might disable the use of <code>LD_LIBRARY_PATH</code>. For more information, see Runtime linking considerations for Linux .</li> </ul> |
| Windows systems        | <p>Set the <code>PATH</code> variable using <code>setmqenv</code>, or <code>crtmqenv</code>.</p>                                                                                                                                                               | <ul style="list-style-type: none"> <li>• No changes required for existing applications.</li> <li>• Easy to change the variable if the location of IBM MQ changes.</li> </ul>                                                        | <ul style="list-style-type: none"> <li>• Environment specific, so must be set in each environment where the application is run.</li> <li>• Possible impact on other applications.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 136. Benefits and drawbacks of the options for loading libraries (continued)

| Platform                         | Option                                                                                                                                                                                                                               | Benefits                                                                                                                                                                                                                                                   | Drawbacks                                                                                                                                                                                                                                                       |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNIX, Linux, and Windows systems | <p>Set the primary installation to a Version 7.1, or later, installation. See “Changing the primary installation” on page 1033.</p> <p>For more information about the primary installation, see Choosing a primary installation.</p> | <ul style="list-style-type: none"> <li>• No changes required for existing applications.</li> <li>• Easy to change the primary installation if the location of IBM MQ changes.</li> <li>• Gives similar behavior to previous versions of IBM MQ.</li> </ul> | <ul style="list-style-type: none"> <li>• When IBM WebSphere MQ Version 7.0.1 is installed, you cannot set the primary installation to Version 7.1, or later.</li> <li>• UNIX and Linux: Does not work if /usr/lib is not in the default search path.</li> </ul> |

## Library loading considerations for HP-UX

The sample compilation commands in the product documentation for previous versions of IBM MQ included the `-W1, +noenvvar` link option for 64-bit applications. This option disables the use of `LD_LIBRARY_PATH` to load shared libraries. If you want your applications to load IBM MQ libraries from a location other than the location specified in the `RPath`, you must update your applications. You can update the applications by recompiling and linking without the `-W1, +noenvvar` link option, or by using the `chatr` command.

To find out how your applications currently load libraries, see “Operating system library loading mechanisms.”

## Library loading considerations for Linux

Applications compiled using some versions of `gcc`, for example, version 3.2.x, can have an embedded `RPath` that cannot be overridden using the `LD_LIBRARY_PATH` environment variable. You can determine if an application is affected by using the `readelf -d applicationName` command. The `RPath` cannot be overridden if the `RPATH` symbol is present and the `RUNPATH` symbol is not present.

## Library loading considerations for Solaris

The sample compilation commands in the product documentation for previous versions of IBM MQ included the `-lmqmcs -lmqmzse` link options. The appropriate versions of these libraries are now loaded automatically by IBM MQ. If IBM MQ is installed in a non-default location, or if there are multiple installations on the system, you must update your applications. You can update the applications by recompiling and linking without the `-lmqmcs -lmqmzse` link options.

## Operating system library loading mechanisms

On Windows systems, several directories are searched to find the libraries:

- The directory the application is loaded from.
- The current directory.
- The directories in the `PATH` environment variable, both the global `PATH` variable and the `PATH` variable of the current user.

On UNIX and Linux systems, there are a number of methods that might have been used to locate the libraries to load:

- Using the `LD_LIBRARY_PATH` environment variable (also `LIBPATH` on AIX, and `SHLIB_PATH` on HP-UX ). If this variable is set, it defines a set of directories that are searched for the required IBM MQ libraries. If any libraries are found in these directories, they are used in preference of any libraries that might be found using the other methods.

- Using an embedded search path (RPath). The application might contain a set of directories to search for the IBM MQ libraries. If the `LD_LIBRARY_PATH` is not set, or if the required libraries were not found using the variable, the RPath is searched for the libraries. If your existing applications use an RPath, but you cannot recompile and link the application, you must either install IBM WebSphere MQ Version 7.1 in the default location, or use another method to find the libraries.
- Using the default library path. If the IBM MQ libraries are not found after searching the `LD_LIBRARY_PATH` variable and RPath locations, the default library path is searched. Usually, this path contains `/usr/lib` or `/usr/lib64`. If the libraries are not found after searching the default library path, the application fails to start because of missing dependencies.

You can use operating system mechanisms to find out if your applications have an embedded search path. For example:

- AIX: `dump`
- HP-UX: `chatr`
- Linux: `readelf`
- Solaris: `elfdump`

#### Related concepts:

“Associating a queue manager with an installation” on page 1034

When you create a queue manager, it is automatically associated with the installation that issued the `crtmqm` command. On UNIX, Linux, and Windows, you can change the installation associated with a queue manager using the `setmqm` command.

“Restrictions for applications using multiple installations”

There are restrictions when using CICS server libraries, fast path connections, message handles, and exits in a multiple installation environment.

“Connecting applications in a multiple installation environment” on page 1024

On UNIX, Linux, and Windows systems, if IBM WebSphere MQ Version 7.1, or later, libraries are loaded, IBM MQ automatically uses the appropriate libraries without you needing to take any further action. IBM MQ uses libraries from the installation associated with the queue manager that the application connects to.

#### Related tasks:

“Changing the primary installation” on page 1033

You can use the `setmqinst` command to set or unset an installation as the primary installation.

#### Related information:

Choosing a primary installation

## Restrictions for applications using multiple installations

There are restrictions when using CICS server libraries, fast path connections, message handles, and exits in a multiple installation environment.

### CICS server libraries

If you are using the CICS server libraries, IBM MQ does not automatically select the correct library level for you. You must compile and link your applications with the appropriate library level for the queue manager to which the application connects. For more information, see *Building libraries for use with TXSeries for Multiplatforms version 5*.

### Message handles

Message handles that use the special value of `MQHC_UNASSOCIATED_HCONN` are limited to use with the first installation loaded in a process. If the message handle cannot be used by a particular installation, reason code `MQRC_HMSG_NOT_AVAILABLE` is returned.

This restriction affects message properties. You cannot use message handles to get message properties from a queue manager on one installation and put them to a queue manager on a different installation. For more information about message handles, see MQCRTMH - Create message handle.

## Exits

In a multiple installation environment, existing exits must be updated for use with IBM WebSphere MQ Version 7.1, or later, installations. Data conversion exits generated using the **crtmqcvx** command must be regenerated using the updated command.

All exits must be written using the MQIEP structure, cannot use an embedded RPATH to locate the IBM MQ libraries, and cannot link to the IBM MQ libraries. For more information, see Writing exits and installable services on UNIX, Linux and Windows .

## Fast path

On a server with multiple installations, applications using a fast path connection to IBM WebSphere MQ Version 7.1 or later must follow these rules:

1. The queue manager must be associated with the same installation as the one from which the application loaded the IBM MQ run time libraries. The application must not use a fast path connection to a queue manager associated with a different installation. An attempt to make the connection results in an error, and reason code MQRC\_INSTALLATION\_MISMATCH.
2. Connecting non-fast path to a queue manager associated with the same installation as the one from which the application has loaded the IBM MQ run time libraries prevents the application connecting fast path, unless either of these conditions are true:
  - The application makes its first connection to a queue manager associated with the same installation a fast path connection.
  - The environment variable, AMQ\_SINGLE\_INSTALLATION is set.
3. Connecting non-fast path to a queue manager associated with a Version 7.1 or later installation, has no effect on whether an application can connect fast path.
4. You cannot combine connecting to a queue manager associated with a Version 7.0.1 installation and connecting fast path to a queue manager associated with a Version 7.1, or later installation.

With AMQ\_SINGLE\_INSTALLATION set, you can make any connection to a queue manager a fast path connection. Otherwise almost the same restrictions apply:

- The installation must be the same one from which the IBM MQ run time libraries were loaded.
- Every connection on the same process must be to the same installation. If you attempt to connect to a queue manager associated with a different installation, the connection fails with reason code MQRC\_INSTALLATION\_MISMATCH. Note that with AMQ\_SINGLE\_INSTALLATION set, this restriction applies to all connections, not only fast path connections.
- Only connect one queue manager with fast path connections.

**Related information:**

MQCONN - Connect queue manager (extended)

MQIEP structure

2583 (0A17) (RC2583): MQRC\_INSTALLATION\_MISMATCH

2587 (0A1B) (RC2587): MQRC\_HMSG\_NOT\_AVAILABLE

2590 (0A1E) (RC2590): MQRC\_FASTPATH\_NOT\_AVAILABLE

**Connecting .NET applications in a multiple installation environment**

By default, applications use the .NET assemblies from the primary installation. If there is no primary installation, or you do not want to use the primary installation assemblies, you must update the application configuration file, or the *DEVPATH* environment variable.

If there is a primary installation on the system, the .NET assemblies and policy files of that installation are registered to the global assembly cache (GAC). The .NET assemblies for all other installations can be found in the installation path of each installation, but the assemblies are not registered to the GAC. Therefore, by default, applications run using the .NET assemblies from the primary installation. You must update the application configuration file if any of the following cases are true:

- You do not have a primary installation.
- You do not want the application to use the primary installation assemblies.
- The primary installation is a lower version of IBM MQ than the version that the application was compiled with.

For information about how to update the application configuration file, see “Connecting .NET applications using the application configuration file.”

You must update the *DEVPATH* environment variable if the following case is true:

- You want your application to use the assemblies from a non-primary installation, but the primary installation is at the same version as the non-primary installation.

For more information about how to update the *DEVPATH* variable, see “Connecting .NET applications using *DEVPATH*” on page 1032.

**Connecting .NET applications using the application configuration file**

Within the application configuration file, you must set various tags to redirect applications to use assemblies that are not from the primary installation.

The following table shows the specific changes that need to be made to the application configuration file to allow .NET applications connect using particular assemblies:

*Table 137. Configuring applications to use particular assemblies*

|                                                                                                                 | Applications compiled with an earlier version of IBM MQ | Applications compiled with a later version of IBM MQ                                                                                                                                                                   |
|-----------------------------------------------------------------------------------------------------------------|---------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| To run an application with a later version IBM MQ primary installation. (later version assemblies in GAC):      | No changes necessary                                    | No changes necessary                                                                                                                                                                                                   |
| To run an application with an earlier version IBM MQ primary installation. (earlier version assemblies in GAC): | No changes necessary                                    | In the application configuration file: <ul style="list-style-type: none"> <li>• Use the <i>&lt;bindingRedirect&gt;</i> tag to indicate the use of the earlier version of the assemblies that are in the GAC</li> </ul> |

Table 137. Configuring applications to use particular assemblies (continued)

|                                                                                                                                        | Applications compiled with an earlier version of IBM MQ                                                                                                                                                                                                                                                       | Applications compiled with a later version of IBM MQ                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| To run an application with a later version of IBM MQ non-primary installation. (later version assemblies in installation folder):      | In the application configuration file: <ul style="list-style-type: none"> <li>• Use the <code>&lt;codebase&gt;</code> tag to point to the location of the later version assemblies</li> <li>• Use the <code>&lt;bindingRedirect&gt;</code> tag to indicate the use of the later version assemblies</li> </ul> | In the application configuration file: <ul style="list-style-type: none"> <li>• Use the <code>&lt;codebase&gt;</code> tag to point to the location of the later version assemblies</li> </ul>                                                                                                                                                                                              |
| To run an application with an earlier version of IBM MQ non-primary installation. (earlier version assemblies in installation folder): | In the application configuration file: <ul style="list-style-type: none"> <li>• Use the <code>&lt;codebase&gt;</code> tag to point to the location of the earlier version assemblies</li> <li>• Include the tag <code>&lt;publisherpolicy Apply=no&gt;</code></li> </ul>                                      | In the application configuration file: <ul style="list-style-type: none"> <li>• Use the <code>&lt;codebase&gt;</code> tag to point to the location of the earlier version assemblies</li> <li>• Use the <code>&lt;bindingRedirect&gt;</code> tag to indicate the use of the earlier version assemblies</li> <li>• Include the tag <code>&lt;publisherpolicy Apply=no&gt;</code></li> </ul> |

A sample application configuration file `NonPrimaryRedirect.config` is shipped in the folder `MQ_INSTALLATION_PATH\tools\dotnet\samples\base`. This file can be modified with the IBM MQ installation path of any non-primary installation. The file can also be directly included in other configuration files using the `<linkedConfiguration>` tag. Samples are provided for `nmqsget.exe.config` and `nmqsput.exe.config`. Both samples use the `<linkedConfiguration>` tag and include the `NonPrimaryRedirect.config` file.

## Connecting .NET applications using DEVPATH

You can find the assemblies using the `DEVPATH` environment variable. The assemblies specified by the `DEVPATH` variable are used in preference to any assemblies in the GAC. See the appropriate Microsoft documentation on `DEVPATH` for more information about when to use this variable.

To find the assemblies using the `DEVPATH` environment variable, you must set the `DEVPATH` variable to the folder that contains the assemblies you want to use. Then, you must then update the application configuration file and add the following runtime configuration information:

```
<configuration>
<runtime>
<developmentMode developerInstallation="true" />
</runtime>
</configuration>
```



### Related concepts:

“Connecting applications in a multiple installation environment” on page 1024

On UNIX, Linux, and Windows systems, if IBM WebSphere MQ Version 7.1, or later, libraries are loaded, IBM MQ automatically uses the appropriate libraries without you needing to take any further action. IBM MQ uses libraries from the installation associated with the queue manager that the application connects to.

### Related information:

Choosing a primary installation

Using .NET

Multiple installations

## Changing the primary installation

You can use the **setmqinst** command to set or unset an installation as the primary installation.

### About this task

This task applies to UNIX, Linux, and Windows.

The primary installation is the installation to which required system-wide locations refer. For more information about the primary installation, and considerations for choosing your primary installation, see [Choosing a primary installation](#).

If an installation of IBM WebSphere MQ Version 7.1 or later is coexisting with an installation of IBM WebSphere MQ Version 7.0.1, the IBM WebSphere MQ Version 7.0.1 installation must be the primary. It is flagged as primary when the IBM WebSphere MQ Version 7.1 or later version is installed, and the IBM WebSphere MQ Version 7.1 or later installation cannot be made primary.

During the installation process on Windows, you can specify that the installation is to be the primary installation. On UNIX and Linux systems, you must issue a **setmqinst** command after installation to set the installation as the primary installation.

“Set the primary installation.”

“Unset the primary installation” on page 1034.

## Set the primary installation

### Procedure

To set an installation as the primary installation:

1. Check if an installation is already the primary installation by entering the following command:

```
MQ_INSTALLATION_PATH/bin/dspmqinst
```

where *MQ\_INSTALLATION\_PATH* is the installation path of a IBM WebSphere MQ Version 7.1 or later installation.

2. If an existing IBM WebSphere MQ Version 7.1 or later installation is set as the primary installation, unset it by following the instructions in “Unset the primary installation” on page 1034. If IBM WebSphere MQ Version 7.0.1 is installed on the system, the primary installation cannot be changed.
3. As root on UNIX and Linux systems, or a member of the Administrators group on Windows systems, enter one of the following commands:

- To set the primary installation using the path of the installation you want to be the primary installation:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -p MQ_INSTALLATION_PATH
```

- To set the primary installation using the name of the installation you want to be the primary installation:

```
MQ_INSTALLATION_PATH/bin/setmqinst -i -n installationName
```

4. On Windows systems, restart the system.

## Unset the primary installation Procedure

To unset an installation as the primary installation:

1. Check which installation is the primary installation by entering the following command:

```
MQ_INSTALLATION_PATH/bin/dspmqinst
```

where *MQ\_INSTALLATION\_PATH* is the installation path of a IBM WebSphere MQ Version 7.1 or later installation.

If IBM WebSphere MQ Version 7.0.1 is the primary installation, you cannot unset the primary installation.

2. As root on UNIX and Linux systems, or a member of the Administrators group on Windows systems, enter one of the following commands:

- To unset the primary installation using the path of the installation you no longer want to be the primary installation:

```
MQ_INSTALLATION_PATH/bin/setmqinst -x -p MQ_INSTALLATION_PATH
```

- To unset the primary installation using the name of the installation you no longer want to be the primary installation:

```
MQ_INSTALLATION_PATH/bin/setmqinst -x -n installationName
```

### Related information:

Features that can be used only with the primary installation on Windows

External library and control command links to primary installation on UNIX and Linux

Uninstalling, upgrading, and maintaining the primary installation

Choosing an installation name

setmqinst

## Associating a queue manager with an installation

When you create a queue manager, it is automatically associated with the installation that issued the **crtmqm** command. On UNIX, Linux, and Windows, you can change the installation associated with a queue manager using the **setmqm** command.

You can use the **setmqm** command in the following ways:

- Moving individual queue managers between equivalent versions of IBM MQ. For example, moving a queue manager from a test to a production system.
- Migrating individual queue managers from an older version of IBM MQ to a newer version of IBM MQ. Migrating queue managers between versions has various implications that you must be aware of. For more information about migrating, see [Migrating and upgrading IBM MQ](#).

To associate a queue manager with an installation:

1. Stop the queue manager using the **endmqm** command from the installation that is currently associated with the queue manager.
2. Associate the queue manager with another installation using the **setmqm** command from that installation.

For example, to set queue manager QMB to be associated with an installation with the name *Installation2*, enter the following command from *Installation2*:

```
MQ_INSTALLATION_PATH/bin/setmqm -m QMB -n Installation2
```

where `MQ_INSTALLATION_PATH` is the path where `Installation2` is installed.

3. Start the queue manager using the **strmqm** command from the installation that is now associated with the queue manager.

This command performs any necessary queue manager migration and results in the queue manager being ready to use.

The installation that a queue manager is associated with limits that queue manager so that it can be administered only by commands from that installation. There are three key exceptions:

- **setmqm** changes the installation associated with the queue manager. This command must be issued from the installation that you want to associate with the queue manager, not the installation that the queue manager is currently associated with. The installation name specified by the **setmqm** command has to match the installation from which the command is issued.
- **strmqm** usually has to be issued from the installation that is associated with the queue manager. However, when a V7.0.1 or earlier queue manager is started on a V7.1 or later installation for the first time, **strmqm** can be used. In this case, **strmqm** starts the queue manager and associates it with the installation from which the command is issued.
- **dspmqr** displays information about all queue managers on a system, not just those queue managers associated with the same installation as the **dspmqr** command. The `dspmqr -o installation` command displays information about which queue managers are associated with which installations.

## Queue manager association in HA environments

For HA environments, the **addmqinf** command automatically associates the queue manager with the installation from which the **addmqinf** command is issued. As long as the **strmqm** command is then issued from the same installation as the **addmqinf** command, no further setup is required. To start the queue manager using a different installation, you must first change the associated installation using the **setmqm** command.

## Queue managers associated with deleted installations

If the installation that a queue manager is associated with has been deleted, or if the queue manager status information is unavailable, the **setmqm** command fails to associate the queue manager with another installation. In this situation, take the following actions:

1. Use the **dspmqrinst** command to see the other installations on your system.
2. Manually modify the `InstallationName` field of the `QueueManager` stanza in `mq.ini` to specify another installation.
3. Use the **dlmqm** command from that installation to delete the queue manager.

## Related concepts:

“Finding installations of IBM MQ on a system”

If you have multiple IBM MQ installations on a system, you can check which versions are installed and where they are.

“The IBM MQ configuration file, mqs.ini” on page 758

The IBM MQ configuration file, mqs.ini, contains information relevant to all the queue managers on the node. It is created automatically during installation.

## Related information:

Choosing a primary installation

setmqm

strmqm

dspmqs

dspmqsinst

## Finding installations of IBM MQ on a system

If you have multiple IBM MQ installations on a system, you can check which versions are installed and where they are.

You can use the following methods to find the IBM MQ installations on your system:

- Use the **dspmqsver** command. This command does not provide details of all installations on a system if it is issued from a Version 7.0.1 installation.
- Use the platform installation tools to query where IBM MQ has been installed. Then use the **dspmqsver** command from a Version 7.1 or later installation. The following commands are examples of commands you can use to query where IBM MQ has been installed:
  - On AIX systems, you can use the **lspp** command:

```
lspp -R ALL -l mqm.base.runtime
```
  - On HP-UX systems, you can use the **swlist** command:

```
swlist -a location -a revision -l product MQSERIES
```
  - On Linux systems, you can use the **rpm** command:

```
rpm -qa --qf "%{NAME}-%{VERSION}-%{RELEASE}\t%{INSTPREFIXES}\n" | grep MQSeriesRuntime
```
  - On Solaris systems, you can use the **pkginfo** and **pkgparam** commands:
    1. List the installed packages by entering the following command:

```
pkginfo | grep -w mqm
```
    2. For each package listed, enter following command:

```
pkgparam pkgname BASEDIR
```
  - On Windows systems, you can use the **wmic** command. This command might install the wmic client:

```
wmic product where "(Name like '%MQ%') AND (not Name like '%bitSupport')"
```

 get Name, Version, InstallLocation
- On UNIX and Linux systems, issue the following command to find out where IBM MQ has been installed:

```
cat /etc/opt/mqm/mqinst.ini
```

Then use the **dspmqsver** command from a Version 7.1 or later installation.

- To display details of installations on the system, on 32-bit Windows, issue the following command:

```
reg.exe query "HKEY_LOCAL_MACHINE\SOFTWARE\IBM\WebSphere MQ\Installation" /s
```
- On 64-bit Windows, issue the following command:

```
reg.exe query "HKEY_LOCAL_MACHINE\SOFTWARE Wow6432Node\IBM\WebSphere MQ\Installation" /s
```

**Note:** the **reg.exe** command will only display information for Version 7.1 or later installations.

### Related information:

dspmqrver

dspmqrinst


Multiple installations


---


## Availability, recovery and restart

Make your applications highly available by maintaining queue availability if a queue manager fails, and recover messages after server or storage failure.

Improve client application availability by using client reconnection to switch a client automatically between a group of queue managers, or to the new active instance of a multi-instance queue manager after a queue manager failure. Automatic client reconnect is not supported by IBM MQ classes for Java.

 Improve server application availability on z/OS by using queue sharing groups.

On Windows,  IBM i, UNIX, and Linux platforms deploy server applications to a multi-instance queue manager, which is configured to run as a single queue manager on multiple servers; if the server running the active instance fails, execution is automatically switched to a standby instance of the same queue manager on a different server. If you configure server applications to run as queue manager services, they are restarted when a standby instance becomes the actively running queue manager instance.

You can configure IBM MQ as part of a platform-specific clustering solution such as Microsoft Cluster Server,  HA clusters on IBM i, or PowerHA® for AIX (formerly HACMP™ on AIX ) and other UNIX and Linux clustering solutions.

Another way to increase server application availability is to deploy server applications to multiple computers in a queue manager cluster.

A messaging system ensures that messages entered into the system are delivered to their destination. IBM MQ can trace the route of a message as it moves from one queue manager to another using the **dspmqrte** command. If a system fails, messages can be recovered in various ways depending on the type of failure, and the way a system is configured.

IBM MQ ensures that messages are not lost by maintaining recovery logs of the activities of the queue managers that handle the receipt, transmission, and delivery of messages. It uses these logs for three types of recovery:

1. *Restart recovery*, when you stop IBM MQ in a planned way.
2. *Failure recovery*, when a failure stops IBM MQ.
3. *Media recovery*, to restore damaged objects.

In all cases, the recovery restores the queue manager to the state it was in when the queue manager stopped, except that any in-flight transactions are rolled back, removing from the queues any updates that were in-flight at the time the queue manager stopped. Recovery restores all persistent messages; nonpersistent messages might be lost during the process.

## Automatic client reconnection

You can make your client applications reconnect automatically, without writing any additional code, by configuring a number of components.

Automatic client reconnection is *inline*. The connection is automatically restored at any point in the client application program, and the handles to open objects are all restored.

In contrast, manual reconnection requires the client application to re-create a connection using MQCONN or MQCONNX, and to reopen objects. Automatic client reconnection is suitable for many, but not all client applications.

Table 138 lists the earliest release of IBM MQ client support that must be installed on a client workstation. You must upgrade client workstations to one of these levels for an application to use automatic client reconnection. Table 139 on page 1039 lists other requirements to enable automatic client reconnection.

With program access to reconnection options, a client application can set reconnection options. Except for JMS and XMS clients, if a client application has access to reconnection options, it can also create an event handler to handle reconnection events.

An existing client application might be able to benefit from reconnection support, without recompilation and linking:

- For a non-JMS client, set the `mqclient.ini` environment variable `DefRecon` to set reconnection options. Use a CCDT to connect to a queue manager. If the client is to connect to a multi-instance queue manager, provide the network addresses of the active and standby queue manager instances in the CCDT.
- For a JMS client, set the reconnection options in the connection factory configuration. When running inside the EJB container of a Java EE server, MDBs can reconnect to IBM MQ using the reconnect mechanism provided by activation specifications of the IBM MQ resource adapter (or listener ports if running in WebSphere Application Server). However, if the application is not an MDB (or is running in the web container) the application must implement its own reconnect logic because automatic client reconnect is not supported in this scenario. The IBM MQ resource adapter provides this reconnect ability for the delivery of messages to message driven beans, but other Java EE elements such as servlets must implement their own reconnection.

Table 138. Supported clients

Client interface	Client	Program access to reconnection options	Reconnection support
Messaging APIs	C, C++, COBOL, Unmanaged Visual Basic, XMS (Unmanaged XMS on Windows)	7.0.1	7.0.1
	JMS (JSE, and Java EE client container and managed containers)	7.0.1.3	7.0.1.3
	IBM MQ classes for Java	Not supported	Not supported
	Managed XMS and managed .NET clients: C#, Visual Basic,	7.1	7.1

Table 138. Supported clients (continued)

Client interface	Client	Program access to reconnection options	Reconnection support
Other APIs	Windows Communication Foundation (Unmanaged <sup>1</sup> )	Not supported	7.0.1
	Windows Communication Foundation (Managed <sup>1</sup> )	Not supported	Not supported
	Axis 1	Not supported	Not supported
	Axis 2	Not supported	7.0.1.3
	HTTP (web 2.0)	Not supported	7.0.1.3

1. Set managed or unmanaged mode in the WCF binding configuration.

Automatic reconnection has the following configuration requirements:

Table 139. Automatic reconnection configuration requirements

Component	Requirement	Effect of not meeting requirement
IBM MQ MQI client installation	See Table 138 on page 1038	MQRC_OPTIONS_ERROR
IBM MQ Server installation	Level 7.0.1	MQRC_OPTIONS_ERROR
Channel	SHARECNV > 0	MQRC_ENVIRONMENT_ERROR
Application environment	Must be threaded	MQRC_ENVIRONMENT_ERROR
MQI	One of: <ul style="list-style-type: none"> <li>• MQCONNX with MQCNO Options set to MQCNO_RECONNECT or MQCNO_RECONNECT_Q_MGR.</li> <li>• Defrecon=YES QMGR in mqclient.ini</li> <li>• In JMS set the CLIENTRECONNECTOPTIONS property of the connection factory.</li> </ul>	MQCC_FAILED when a connection is broken or queue manager ends or fails.

Figure 150 on page 1040 shows the main interactions between components that are involved in client reconnection.

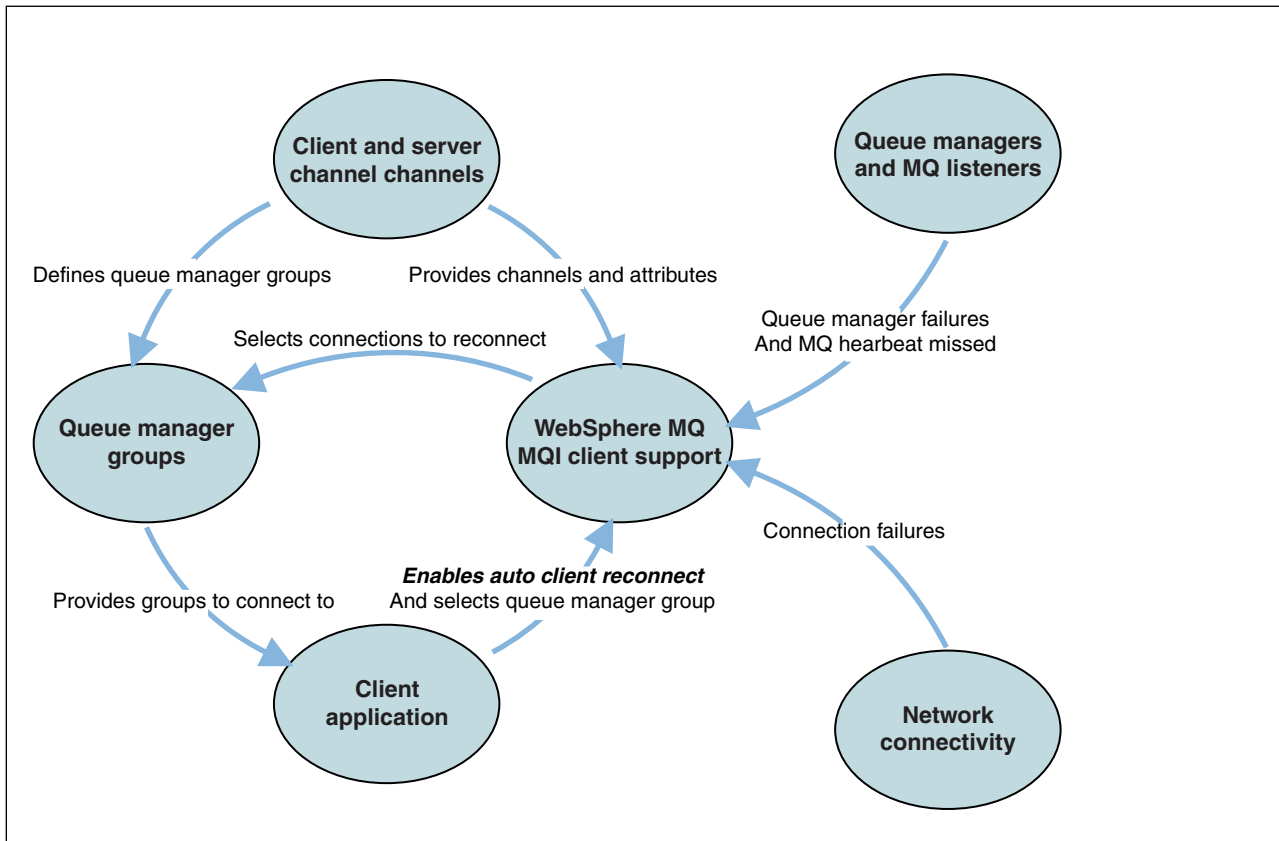


Figure 150. Automatic client reconnection

## Client application

The client application is an IBM MQ MQI client.

- By default clients are not automatically reconnected. Enable the automatic client reconnection by setting the MQCONNX MQCNO Option MQCNO\_RECONNECT or MQCNO\_RECONNECT\_Q\_MGR.
- Many applications are written in such a way that they are able to take advantage of auto-reconnection with no additional coding. Enable automatic reconnection for existing programs, without making any coding changes, by setting the DefRecon attribute in the channels stanza of the mqclient.ini configuration file.
- Use one of these three options:
  1. Modify the program so that the logic is unaffected by reconnection. For example, you might have to issue MQI calls within the sync point, and resubmit backed-out transactions.
  2. Add an event handler to detect reconnection, and restore the state of the client application when the connection is reestablished.
  3. Do not enable auto-reconnection: instead, disconnect the client and issue a new MQCONN or MQCONNX MQI call to find another queue manager instance that is running in the same queue manager group.

For further details about these three options, see “Application recovery” on page 1123.

- Reconnecting to a queue manager of the same name does not guarantee that you have reconnected to the same instance of a queue manager.

Use an MQCNO option MQCNO\_RECONNECT\_Q\_MGR, to reconnect to an instance of the same queue manager.

- A client can register an event handler so that it can be informed the state of reconnection. The MQHCONN passed in the event handler cannot be used. The following reason codes are provided:



### **MQRC\_RECONNECTING**

The connection failed, and the system is attempting to reconnect. You receive multiple MQRC\_RECONNECTING events if multiple reconnect attempts are made.

### **MQRC\_RECONNECTED**

The reconnection made and all handles successfully reestablished.

### **MQRC\_RECONNECT\_FAILED**

The reconnection was not successful.

### **MQRC\_RECONNECT\_QMID\_MISMATCH**

A reconnectable connection specified MQCNO\_RECONNECT\_Q\_MGR and the connection attempted to reconnect to a different queue manager.

### **MQRC\_RECONNECT\_Q\_MGR\_REQD**

An option, such MQMO\_MATCH\_MSG\_TOKEN in an MQGET call, was specified in the client program that requires reconnection to the same queue manager.

- A reconnectable client is able to reconnect automatically only *after* connecting. That is, the MQCONN call itself is not tried again if it fails. For example, if you receive the return code 2543 - MQRC\_STANDBY\_Q\_MGR from MQCONN, reissue the call after a short delay.

### **MQRC\_RECONNECT\_INCOMPATIBLE**

This reason code is returned when the application tries to use MQPMO\_LOGICAL\_ORDER (with MQPUT and MQPUT1) or MQGMO\_LOGICAL\_ORDER (with MQGET ) when reconnect options are set. The reason for returning the reason code is to make sure that applications never use reconnect in such cases.

### **MQRC\_CALL\_INTERRUPTED**

This reason code is returned when the connection breaks during the execution of Commit call and the client reconnects. An MQPUT of a persistent message outside the sync point also results in the same reason code being returned to the application.

## **Multi-instance queue managers**

Simplify restarting IBM MQ MQI client applications, after a multi-instance queue manager has activated its standby instance, by using automatic client reconnection.

The standby instance of a multi-instance queue manager is typically at a different network address to the active instance. Include the network addresses of both the instances in the client connection definition table (CCDT). Either provide a list of network addresses for the **CONNNAME** parameter, or define multiple rows for the queue manager in the CCDT.

Commonly, IBM MQ MQI clients reconnect to any queue manager in a queue manager group. Sometimes you want an IBM MQ MQI client to reconnect only to the same queue manager. It might have an affinity to a queue manager. You can prevent a client from reconnecting to a different queue manager. Set the MQCNO option, MQCNO\_RECONNECT\_Q\_MGR. The IBM MQ MQI client fails if it reconnects to a different queue manager. If you set the MQCNO option, MQCNO\_RECONNECT\_Q\_MGR, do not include other queue managers in the same queue manager group. The client returns an error if the queue manager it reconnects to is not the same queue manager as the one it connected to.

## **Queue manager groups**

You can select whether the client application always connects and reconnects to a queue manager of the same name, to the same queue manager, or to any of a set of queue managers that are defined with the same QMNAME value in the client connection table.

- The queue manager *name* attribute, QMNAME, in the client channel definition is the name of a queue manager group.

- In your client application, if you set the value of the MQCONN or MQCONNX QmgrName parameter to a queue manager name, the client connects only to queue managers with that name. If you prefix the queue manager name with an asterisk(\*), the client connects to any queue manager in the queue manager group with the same QMNAME value. For a full explanation, see Queue manager groups in the CCDT.

## Queue sharing groups

Automatic client reconnection to z/OS queue sharing groups, uses the same mechanisms for reconnection as any other environment. The client will reconnect to the same selection of queue managers as is configured for the original connection. For example, when using the client channel definition table the administrator should ensure that all entries in the table, resolve to the same z/OS queue sharing group.

## Client and server channel definitions

Client and server channel definitions define the groups of queue managers a client application can reconnect to. The definitions govern the selection and timing of reconnections, and other factors, such as security; see the related topics. The most relevant channel attributes to consider for reconnection are listed in two groups:

### Client connection attributes

#### Connection affinity (AFFINITY) AFFINITY

Connection affinity.

#### Client channel weight (CLNTWGHT) CLNTWGHT

Client channel weight.

#### Connection name (CONNAME) CONNAME

Connection information.

#### Heartbeat interval (HBINT) HBINT

Heartbeat interval. Set the heartbeat interval on the server connection channel.

#### Keepalive Interval (KAINT) KAINT

Keepalive interval. Set the keepalive interval on the server connection channel.

Note that KAINT applies to z/OS only.

#### Queue manager name (QMNAME) QMNAME

Queue manager name.

### Server connection attributes

#### Heartbeat interval (HBINT) HBINT

Heartbeat interval. Set the heartbeat interval on the client connection channel.

#### Keepalive Interval (KAINT) KAINT

Keepalive interval. Set the keepalive interval on the client connection channel.

Note that KAINT applies to z/OS only.

KAINT is a network layer heartbeat, and HBINT is an IBM MQ heartbeat between the client and the queue manager. Setting these heartbeats to a shorter time serves two purposes:

1. By simulating activity on the connection, network layer software that is responsible for closing inactive connections is less likely to shut down your connection.
2. If the connection is shut down, the delay before the broken connection is detected, is shortened.

The default TCP/IP keepalive interval is two hours. Consider setting the KAINT and HBINT attributes to a shorter time. Do not assume that the normal behavior of a network suits the needs of automatic

reconnection. For example, some firewalls can shut down an inactive TCP/IP connection after as little as 10 minutes.

## Network connectivity

Only network failures that are passed to the IBM MQ MQI client by the network, are handled by the automatic reconnection capability of the client.

- Reconnections performed automatically by the transport are invisible to IBM MQ.
- Setting HBINT helps to deal with network failures that are invisible to IBM MQ.

## Queue managers and IBM MQ listeners

Client reconnection is triggered by server failure, queue manager failure, network connectivity failure, and by an administrator switching over to another queue manager instance.

- If you are using a multi-instance queue manager, an additional cause of client reconnection occurs when you switch control from the active queue manager instance to a standby instance.
- Ending a queue manager using the default **endmqm** command, does not trigger automatic client reconnection. Add the **-r** option on the **endmqm** command to request automatic client reconnection, or the **-s** option to transfer to a standby queue manager instance after shutting down.

## IBM MQ MQI client automatic reconnection support

If you use the automatic client reconnection support in the IBM MQ MQI client, the client application automatically reconnects and continues processing without you issuing an MQCONN or MQCONNX MQI call to reconnect to the queue manager.

- Automatic client reconnection is triggered by one of the following occurrences:
  - queue manager failure
  - ending a queue manager and specifying the **-r**, reconnect, option on the **endmqm** command
- The MQCONNX MQCNO options control whether you have enabled the automatic client reconnection. The options are described in Reconnection options.
- Automatic client reconnection issues MQI calls on behalf of your application to restore the connection handle and the handles to other open objects, so that your program can resume normal processing after it has processed any MQI errors that resulted from the broken connection. See “Recovery of an automatically reconnected client” on page 1125.
- If you have written a channel exit program for the connection, the exit receives these additional MQI calls.
- You can register a reconnection event handler, which is triggered when reconnection begins and when it finishes.

Although the intended reconnection time is no more than a minute, reconnection can take longer because a queue manager might have numerous resources to manage. During this time, a client application might be holding locks that do not belong to IBM MQ resources. There is a timeout value you can configure to limit the time a client waits for reconnection. The value (in seconds) is set in the `mqclient.ini` file.

```
Channels:
MQReconnectTimeout = 1800
```

No reconnection attempts are made after the timeout has expired. When the system detects that the timeout has expired it returns a MQRC\_RECONNECT\_FAILED error.

## Console message monitoring

z/OS

There are a number of information messages issued by the queue manager or channel initiator that should be considered particularly significant. These messages do not in themselves indicate a problem, but may be useful in tracking because they do indicate a potential issue which might need addressing.

The presence of these console messages might also indicate that a user application is putting a large number of messages to the page set, which might be a symptom of a larger problem:

- A problem with the user application which PUTs messages, such as an uncontrolled loop.
- A user application which GETs the messages from the queue is no longer functioning.

distributed

### Cluster error recovery for servers on distributed platforms

From IBM WebSphere MQ Version 7.1 onwards, the queue manager reruns operations that caused problems, until the problems are resolved. See Changes to cluster error recovery on servers other than z/OS for more information.

z/OS

### Console messages to monitor

The following list outlines messages which can potentially indicate larger problems. Determine if it is necessary to track these messages with system automation and provide appropriate documentation so any potential problems can be followed up effectively.

**CSQI004I:** *csect-name* **CONSIDER INDEXING** *queue-name* **BY** *index-type* **FOR** *connection-type* **CONNECTION** *connection-name, num-msgs* **MESSAGES SKIPPED**

- The queue manager has detected an application receiving messages by message ID or correlation ID from a queue that does not have an index defined.
- Consider establishing an index for the identified queue by altering the local queue object, *queue-name*, INDXTYPE attribute to have value *index-type*.

**CSQI031I:** *csect-name* **THE NEW EXTENT OF PAGE SET** *psid* **HAS FORMATTED SUCCESSFULLY**

- Check the curdepth of the queues allocated to this page set.
- Investigate the cause of the failure to process the messages.

**CSQI041I:** *csect-name* **JOB** *jobname* **USER** *userid* **HAD ERROR ACCESSING PAGE SET** *psid*

- Determine if the page set is allocated to the queue manager.
- Issue a **DISPLAY USAGE** command to determine the state of the page set.
- Check the queue manager joblog for additional error messages.

**CSQJ004I:** **ACTIVE LOG COPY** *n* **INACTIVE, LOG IN SINGLE MODE, ENDRBA=** *ttt*

- The queue manager has activated 'single' logging mode. This is often indicative of a log offload problem.
- Issue a **DISPLAY LOG** command to determine your settings for duplexing of active and archive logs. This display also shows how many active logs need offload processing.
- Check the queue manager joblog for additional error messages

**CSQJ114I:** **ERROR ON ARCHIVE DATA SET, OFFLOAD CONTINUING WITH ONLY ONE ARCHIVE DATA SET BEING GENERATED**

- Check the queue manager joblog for additional error messages.
- Make a second copy of the archive log and update your BSDS manually.

**CSQJ004I: ACTIVE LOG COPY *n* INACTIVE, LOG IN SINGLE MODE, ENDRBA= *ttt***

- Allocate enough units to enable archiving.
- Check the queue manager joblog for additional error messages.

**CSQJ136I: UNABLE TO ALLOCATE TAPE UNIT FOR CONNECTION-ID= *xxxx* CORRELATION-ID= *yyyyyy*, *m* ALLOCATED *n* ALLOWED**

- Check the queue manager joblog for additional error messages.

**CSQJ151I: *csect-name* ERROR READING RBA *rrr*, CONNECTION-ID= *xxxx* CORRELATION-ID= *yyyyyy* REASON CODE= *ccc***

- Check the queue manager joblog for additional messages.
- Issue a **DISPLAY CONN** command to determine which connection is not committing its activity.
- Ensure the application can commit its updates.

**CSQJ160I: LONG-RUNNING UOW FOUND, URID= *urid* CONNECTION NAME= *name***

- Check the queue manager joblog for additional messages.
- Issue a **DISPLAY CONN** command to determine which connection is not committing its activity.
- Ensure the application can commit its updates.

**CSQJ161I: UOW UNRESOLVED AFTER *n* OFFLOADS, URID= *urid* CONNECTION NAME= *name***

- Determine if the page set is allocated to the queue manager.
- Issue a **DISPLAY USAGE** command to determine the state of the page set.
- Check the queue manager joblog for additional messages.

**CSQP011E: CONNECT ERROR STATUS *ret-code* FOR PAGE SET *psid***

- Check the curdepth of the queues allocated to this page set.
- Investigate the cause of the failure to process messages.

**CSQP013I: *csect-name* NEW EXTENT CREATED FOR PAGE SET *psid*. NEW EXTENT WILL NOW BE FORMATTED**

- Check the curdepth of the queues allocated to this page set.
- Investigate the cause of failure to process messages.
- Determine if queues need to be relocated to another page set.
- If the volume is full, determine if you need to make the page set a multi volume dataset. If the page set is already multi-volume, consider adding more volumes to the storage group being used. Once more space is available retry the expansion by setting the page set **EXPAND** method to **SYSTEM**. If a retry is required, toggle **EXPAND** to **SYSTEM** and then back to your normal setting.

**CSQP014E: *csect-name* EXPANSION FAILED FOR PAGE SET *psid*. FUTURE REQUESTS TO EXTEND IT WILL BE REJECTED**

- Check the curdepth of the queues allocated to this page set.
- Investigate the cause of failure to process messages.
- Determine if queues need to be relocated to another page set.

**CSQP016E: *csect-name* PAGE SET *psid* HAS REACHED THE MAXIMUM NUMBER OF EXTENTS. IT CANNOT BE EXTENDED AGAIN**

- Check the curdepth of the queues allocated to this page set.
- Investigate the cause of failure to process messages.

- CSQP017I:** *csect-name* **EXPANSION STARTED FOR PAGE SET** *psid*
- Check the IMS log to determine why the Units of Work are still indoubt.
  - Issue DISPLAY THREAD commands to determine the state of the Units of Work in IBM MQ.
- CSQP047E:** **Unavailable page sets can cause problems - take action to correct this situation**
- Follow the System Programmer Response.
- CSQQ008I:** *nn* **units of recovery are still in doubt in queue manager** *qqqq*
- Investigate the state of your dead letter queue. Ensure the dead letter queue is not PUT disabled.
  - Ensure the dead letter queue is not at the MAXMSG limit.
- CSQQ113I:** *psb-name region-id* **This message cannot be processed**
- Check the CSQOUTX dataset to determine the cause of the CSQINPX failure.
  - Some commands may not be processed.
- CSQX035I:** *csect-name* **Connection to queue manager** *qmgr-name* **stopping or broken, MQCC=** *mqqc* **MQRC=** *mqrq (mqrq-text)*
- Check the MQRC to determine the cause of the failure.
  - These codes are documented in IBM MQ for z/OS messages, completion, and reason codes.
- CSQX032I:** *csect-name* **Initialization command handler terminated**
- Check the MQRC to determine the cause of the failure.
  - These codes are documented in IBM MQ for z/OS messages, completion, and reason codes.
- CSQX048I:** *csect-name* **Unable to convert message for** *name*, **MQCC=** *mqqc* **MQRC=** *mqrq (mqrq-text)*
- Check the joblog to determine the cause of the TCP/IP failure.
  - Check the TCP/IP address space for errors.
- CSQX234I:** *csect-name* **Listener stopped, TRPTYPE=** *trptype* **INDISP=** *disposition*
- If the listener does not stop, following a **STOP** command, check the TCP/IP address space for errors.
  - Follow the Systems Programmer Response.
- CSQX407I:** *csect-name* **Cluster queue** *q-name* **definitions inconsistent**
- Multiple cluster queues within the cluster have inconsistent values. Investigate and resolve the differences.
- CSQX411I:** *csect-name* **Repository manager stopped**
- If the repository manager has stopped because of an error, check the joblog for messages.
- CSQX417I:** *csect-name* **Cluster-senders remain for removed queue manager** *qmgr-name*
- Follow the System Programmer Response.
- CSQX418I:** *csect-name* **Only one repository for cluster** *cluster-name*
- For increased high availability, clusters should be configured with two full repositories.
- CSQX419I:** *csect-name* **No cluster-receivers for cluster** *cluster-name*
- Follow the System Programmer Response.
- CSQX420I:** *csect-name* **No repositories for cluster** *cluster-name*
- Follow the System Programmer Response.
- CSQX448E:** *csect-name* **Repository manager stopping because of errors. Restart in** *n* **seconds**
- Follow the System Programmer Response.
- This message is put out every 600 seconds (10 minutes) until the SYSTEM.CLUSTER.COMMAND.QUEUE is enabled, by using the command:

```
ALTER QLOCAL(SYSTEM.CLUSTER.COMMAND.QUEUE) GET(ENABLED)
```

Before enabling the queue, manual intervention might be required to resolve the problem that caused the repository manager to end, prior to the first CSQX448E message being issued.

## Using IBM MQ with high availability configurations

If you want to operate your IBM MQ queue managers in a high availability (HA) configuration, you can set up your queue managers to work either with a high availability manager, such as PowerHA for AIX (formerly HACMP ) or the Microsoft Cluster Service (MSCS), or with IBM MQ multi-instance queue managers.

You need to be aware of the following configuration definitions:

### Queue manager clusters

Groups of two or more queue managers on one or more computers, providing automatic interconnection, and allowing queues to be shared among them for load balancing and redundancy.

### HA clusters

HA clusters are groups of two or more computers and resources such as disks and networks, connected together and configured in such a way that, if one fails, a high availability manager, such as HACMP ( UNIX ) or MSCS ( Windows ) performs a *failover*. The failover transfers the state data of applications from the failing computer to another computer in the cluster and re-initiates their operation there. This provides high availability of services running within the HA cluster. The relationship between IBM MQ clusters and HA clusters is described in “Relationship of HA clusters to queue manager clusters” on page 1048.

### Multi-instance queue managers

Instances of the same queue manager configured on two or more computers. By starting multiple instances, one instance becomes the active instance and the other instances become standbys. If the active instance fails, a standby instance running on a different computer automatically takes over. You can use multi-instance queue managers to configure your own highly available messaging systems based on IBM MQ, without requiring a cluster technology such as HACMP or MSCS. HA clusters and multi-instance queue managers are alternative ways of making queue managers highly available. Do not combine them by putting a multi-instance queue manager in an HA cluster.

## Differences between multi-instance queue managers and HA clusters

Multi-instance queue managers and HA clusters are alternative ways to achieve high availability for your queue managers. Here are some points that highlight the differences between the two approaches.

Multi-instance queue managers include the following features:

- Basic failover support integrated into IBM MQ
- Faster failover than HA cluster
- Simple configuration and operation
- Integration with MQ Explorer

Limitations of multi-instance queue managers include:

- Highly available, high performance networked storage required
- More complex network configuration because queue manager changes IP address when it fails over

HA clusters include the following features:

- The ability to coordinate multiple resources, such as an application server or database
- More flexible configuration options including clusters comprising more than two nodes

- Can failover multiple times without operator intervention
- Takeover of queue manager's IP address as part of the failover

Limitations of HA clusters include:

- Additional product purchase and skills are required
- Disks which can be switched between the nodes of the cluster are required
- Configuration of HA clusters is relatively complex
- Failover is rather slow historically, but recent HA cluster products are improving this
- Unnecessary failovers can occur if there are shortcomings in the scripts that are used to monitor resources such as queue managers

## Relationship of HA clusters to queue manager clusters

Queue manager clusters provide load balancing of messages across available instances of queue manager cluster queues. This offers higher availability than a single queue manager because, following a failure of a queue manager, messaging applications can still send messages to, and access, surviving instances of a queue manager cluster queue. However, although queue manager clusters automatically route new messages to the available queue managers in a cluster, messages currently queued on an unavailable queue manager are not available until that queue manager is restarted. For this reason, queue manager clusters alone do not provide high availability of all message data or provide automatic detection of queue manager failure and automatic triggering of queue manager restart or failover. High Availability (HA) clusters provide these features. The two types of cluster can be used together to good effect. For an introduction to queue manager clusters, see *Designing clusters*.

### Using IBM MQ with a high availability cluster on UNIX and Linux

You can use IBM MQ with a high availability (HA) cluster on UNIX and Linux platforms: for example, PowerHA for AIX (formerly HACMP), Veritas Cluster Server, HP Serviceguard, or a Red Hat Enterprise Linux cluster with Red Hat Cluster Suite.

Before IBM WebSphere MQ Version 7.0.1, SupportPac MC91 was provided to assist in configuring HA clusters. IBM WebSphere MQ Version 7.0.1 provided a greater degree of control than previous versions over where queue managers store their data. This makes it easier to configure queue managers in an HA cluster. Most of the scripts provided with SupportPac MC91 are no longer required, and the SupportPac is withdrawn.

This section introduces “HA cluster configurations,” the relationship of HA clusters to queue manager clusters, “IBM MQ clients” on page 1049, and “IBM MQ operating in an HA cluster” on page 1049, and guides you through the steps and provides example scripts that you can adapt to configure queue managers with an HA cluster.

Refer to the HA cluster documentation particular to your environment for assistance with the configuration steps described in this section.

### HA cluster configurations

In this section the term *node* is used to refer to the entity that is running an operating system and the HA software; “computer”, “system” or “machine” or “partition” or “blade” might be considered synonyms in this usage. You can use IBM MQ to help set up either standby or takeover configurations, including mutual takeover where all cluster nodes are running IBM MQ workload.

A *standby* configuration is the most basic HA cluster configuration in which one node performs work while the other node acts only as standby. The standby node does not perform work and is referred to as idle; this configuration is sometimes called *cold standby*. Such a configuration requires a high degree of hardware redundancy. To economize on hardware, it is possible to extend this configuration to have



multiple worker nodes with a single standby node. The point of this is that the standby node can take over the work of any other worker node. This configuration is still referred to as a standby configuration and sometimes as an "N+1" configuration.

A *takeover* configuration is a more advanced configuration in which all nodes perform some work and critical work can be taken over in the event of a node failure.

A *one-sided takeover* configuration is one in which a standby node performs some additional, noncritical and unmovable work. This configuration is similar to a standby configuration but with (noncritical) work being performed by the standby node.

A *mutual takeover* configuration is one in which all nodes are performing highly available (movable) work. This type of HA cluster configuration is also sometimes referred to as "Active/Active" to indicate that all nodes are actively processing critical workload.

With the extended standby configuration or either of the takeover configurations it is important to consider the peak load that might be placed on a node that can take over the work of other nodes. Such a node must possess sufficient capacity to maintain an acceptable level of performance.

## Relationship of HA clusters to queue manager clusters

Queue manager clusters reduce administration and provide load balancing of messages across instances of queue manager cluster queues. They also offer higher availability than a single queue manager because, following a failure of a queue manager, messaging applications can still access surviving instances of a queue manager cluster queue. However, queue manager clusters alone do not provide automatic detection of queue manager failure and automatic triggering of queue manager restart or failover. HA clusters provide these features. The two types of cluster can be used together to good effect.

## IBM MQ clients

IBM MQ clients that are communicating with a queue manager that might be subject to a restart or takeover must be written to tolerate a broken connection and must repeatedly attempt to reconnect. IBM WebSphere MQ Version 7 introduced features in the processing of the Client Channel Definition Table (CCDT) that assist with connection availability and workload balancing; however these are not directly relevant when working with a failover system.

Transactional functionality allows an IBM MQ MQI client to participate in two-phase transactions, as long as the client is connected to the same queue manager. Transactional functionality cannot use techniques, such as an IP load balancer, to select from a list of queue managers. When you use an HA product, a queue manager maintains its identity (name and address) whichever node it is running on, so transactional functionality can be used with queue managers that are under HA control.

## IBM MQ operating in an HA cluster

All HA clusters have the concept of a unit of failover. This is a set of definitions that contains all the resources that make up the highly available service. The unit of failover includes the service itself and all other resources upon which it depends.

HA solutions use different terms for a unit of failover:

- On PowerHA for AIX the unit of failover is called a *resource group*.
- On Veritas Cluster Server it is known as a *service group*.
- On Serviceguard it is called a *package*.

This topic uses the term *resource group* to mean a unit of failover.

The smallest unit of failover for IBM MQ is a queue manager. Typically, the resource group containing the queue manager also contains shared disks in a volume group or disk group that is reserved exclusively for use by the resource group, and the IP address that is used to connect to the queue manager. It is also possible to include other IBM MQ resources, such as a listener or a trigger monitor in the same resource group, either as separate resources, or under the control of the queue manager itself.

A queue manager that is to be used in an HA cluster must have its data and logs on disks that are shared between the nodes in the cluster. The HA cluster ensures that only one node in the cluster at a time can write to the disks. The HA cluster can use a monitor script to monitor the state of the queue manager.

It is possible to use a single shared disk for both the data and logs that are related to the queue manager. However, it is normal practice to use separate shared file systems so that they can be independently sized and tuned.

Figure 1 illustrates a HA cluster with two nodes. The HA cluster is managing the availability of a queue

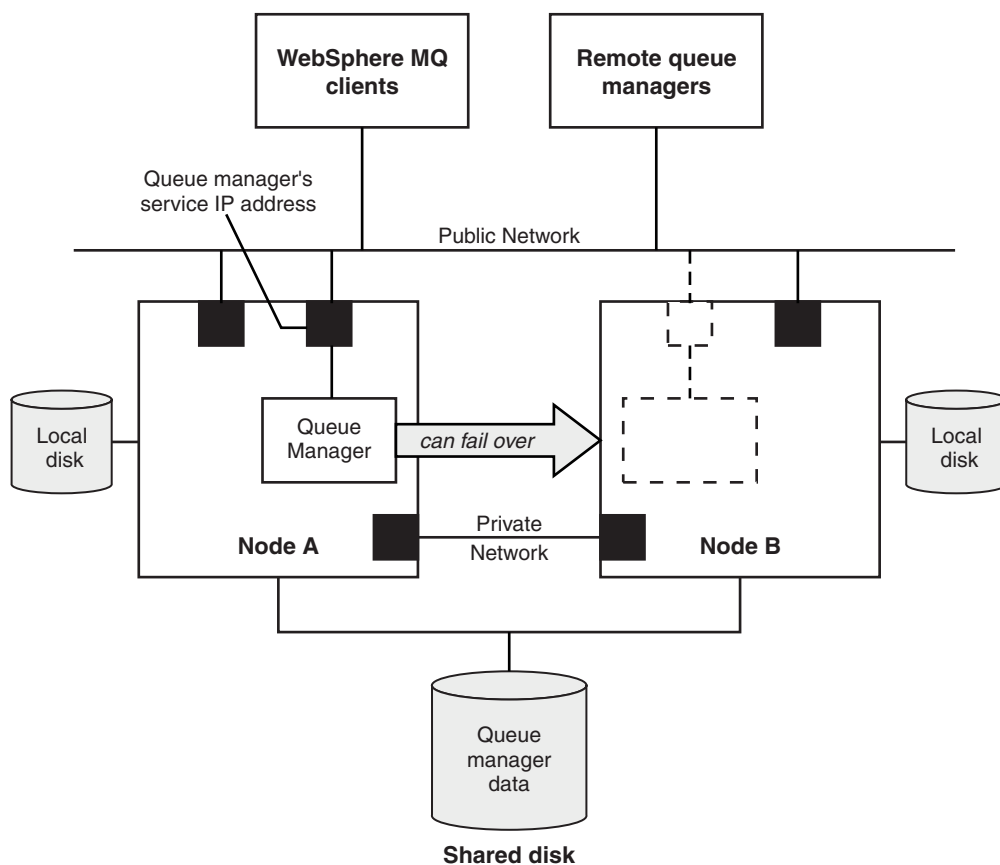


Figure 151. HA cluster

manager which has been defined in a resource group. This is an active/passive or cold standby configuration, because only one node, node A, is currently running a queue manager. The queue manager was created with its data and log files on a shared disk. The queue manager has a service IP address which is also managed by the HA cluster. The queue manager depends on the shared disk and its service IP address. When the HA cluster fails the queue manager over from node A to node B, it first moves the queue manager's dependent resources onto node B and then starts the queue manager.

If the HA cluster contains more than one queue manager, your HA cluster configuration might result in two or more queue managers running on the same node after a failover. Each queue manager in the HA cluster must be assigned its own port number, which it uses on whichever cluster node it happens to be active at any particular time.

Generally, the HA cluster runs as the root user. IBM MQ runs as the mqm user. Administration of IBM MQ is granted to members of the mqm group. Ensure that the mqm user and group both exist on all HA cluster nodes. The user ID and group ID must be consistent across the cluster. Administration of IBM MQ by the root user is not allowed; scripts that start, stop, or monitor scripts must switch to the mqm user.

**Note:** IBM MQ must be installed correctly on all nodes; you cannot share the product executable files.

### Configuring the shared disks:

An IBM MQ queue manager in an HA cluster requires data files and log files to be in common named remote file systems on a shared disk.

To configure the shared disks, complete the following steps:

1. Decide the names of the mount points for the queue manager's file systems. For example, /MQHA/qmgrname/data for the queue manager's data files and /MQHA/qmgrname/log for its log files.
2. Create a volume group (or disk group) to contain the queue manager's data and log files. This volume group is managed by the high availability (HA) cluster in the same resource group as the queue manager.
3. Create the file systems for the queue manager's data and log files in the volume group.
4. For each node in turn, create the mount points for the file systems and make sure that the file systems can be mounted. The mqm user must own the mount points.

Figure 1 shows a possible layout for a queue manager in an HA cluster. The queue manager's data and log directories are both on the shared disk which is mounted at /MQHA/QM1. This disk is switched between the nodes of the HA cluster when failover occurs so that the data is available wherever the queue manager is restarted. The mqs.ini file has a stanza for the QM1 queue manager. The Log stanza in the qm.ini file has a value for LogPath.

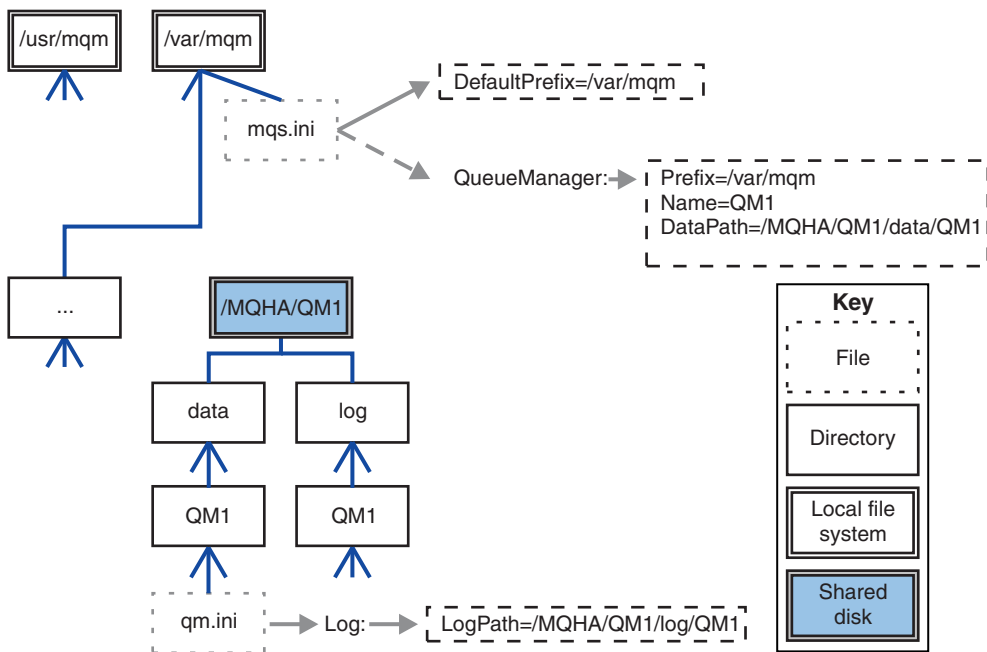


Figure 152. Shared named data and log directories

## Creating a queue manager for use in a high availability (HA) cluster:

The first step towards using a queue manager in a high availability cluster is to create the queue manager on one of the nodes.

To create a queue manager for use in an HA cluster, select one of the nodes in the cluster on which to create the queue manager. On this node complete the following steps:

1. Mount the queue manager's file systems on the node.
2. Create the queue manager by using the **crtmqm** command. For example:  

```
crtmqm -md /MQHA/qmgrname/data -ld /MQHA/qmgrname/log qmgrname
```
3. Start the queue manager manually by using the **strmqm** command.
4. Complete any initial configuration of the queue manager, such as creating queues and channels, and setting the queue manager to start a listener automatically when the queue manager starts.
5. Stop the queue manager by using the **endmqm** command.
6. Use the **dspmqlnf** command to display the **addmqinf** command that you can use in a later task, which is documented in "Adding queue manager configuration information to other nodes in a high availability (HA) cluster":

```
dspmqlnf -o command qmgrname
```

where `qmgrname` is the name of the queue manager.

7. The **addmqinf** command that is displayed will be similar to the following example:

```
addmqinf -sQueueManager -vName=qmgrname -vDirectory=qmgrname \
-vPrefix=/var/mqm -vDataPath=/MQHA/qmgrname/data/qmgrname
```

Make a careful note of the displayed command.

8. Unmount the queue manager's file systems.

You are now ready to complete the steps described in "Adding queue manager configuration information to other nodes in a high availability (HA) cluster."

## Adding queue manager configuration information to other nodes in a high availability (HA) cluster:

You must add the queue manager configuration to the other nodes in the HA cluster.

Before you complete this task, you must have completed the steps in "Creating a queue manager for use in a high availability (HA) cluster."

To add the configuration information for the queue manager to each of other nodes in the HA cluster, complete the following steps on each of the other nodes:

1. Mount the queue manager file systems.
2. Add the queue manager configuration information to the node, either by editing `/var/mqm/mqs.ini` directly, or by issuing the **addmqinf** command that was displayed by the **dspmqlnf** command in steps 6 and 7 in "Creating a queue manager for use in a high availability (HA) cluster."
3. Start and stop the queue manager to verify the configuration.

The commands used to start and stop the queue manager must be issued from the same IBM MQ installation as the **addmqinf** command. To start and stop the queue manager from a different installation, you must first set the installation associated with the queue manager using the **setmqm** command. For more information, see `setmqm`.

4. Unmount the queue manager file systems.

## Starting a queue manager under control of a high availability (HA) cluster:

The queue manager is represented in the HA cluster as a resource. The HA cluster must be able to start and stop the queue manager. In most cases you can use a shell script to start the queue manager. You must make these scripts available at the same location on all nodes in the cluster, either using a network filesystem or by copying them to each of the local disks.

**Note:** Before you restart a failed queue manager, you must disconnect your applications from that instance of the queue manager. If you do not, the queue manager might not restart correctly.

Examples of suitable shell scripts are given here. You can tailor these to your needs and use them to start the queue manager under the control of your HA cluster.

The following shell script is an example of how to switch from the HA cluster user to the mqm user so that the queue manager can be successfully started:

```
#!/bin/ksh
A simple wrapper script to switch to the mqm user.
su mqm -c name_of_your_script $*
```

The following shell script is an example of how to start a queue manager without making any assumptions about the current state of the queue manager. Note that it uses an extremely abrupt method of ending any processes that belong to the queue manager:

```
#!/bin/ksh
#
This script robustly starts the queue manager.
#
The script must be run by the mqm user.
#
The only argument is the queue manager name. Save it as QM variable
QM=$1

if [-z "$QM"]
then
 echo "ERROR! No queue manager name supplied"
 exit 1
fi

End any queue manager processes which might be running.

srchstr="(|-m)$QM *.*$"
for process in amqzmuc0 amqzma0 amqfcxba amqfqpub amqpcsea amqzlaa0 \
 amqzlsa0 runmqchi runmqlsr amqcrsta amqrrmfa amqrmppa \
 amqzfuma amqzmuf0 amqzmur0 amqzmgr0
do
 ps -ef | tr "\t" " " | grep $process | grep -v grep | \
 egrep "$srchstr" | awk '{print $2}' | \
 xargs kill -9 > /dev/null 2>&1
done

It is now safe to start the queue manager.
The strmqm command does not use the -x flag.
strmqm ${QM}
```

You can modify the script to start other related programs.

## Stopping a queue manager under the control of a high availability (HA) cluster:

In most cases, you can use a shell script to stop a queue manager. Examples of suitable shell scripts are given here. You can tailor these to your needs and use them to stop the queue manager under control of your HA cluster.

The following script is an example of how to immediately stop without making assumptions about the current state of the queue manager. The script must be run by the mqm user; it might therefore be necessary to wrap this script in a shell script to switch the user from the HA cluster user to mqm (an example shell script is provided in "Starting a queue manager under control of a high availability (HA) cluster" on page 1053):

```
#!/bin/ksh
#
The script ends the QM by using two phases, initially trying an immediate
end with a time-out and escalating to a forced stop of remaining
processes.
#
The script must be run by the mqm user.
#
There are two arguments: the queue manager name and a timeout value.
QM=$1
TIMEOUT=$2

if [-z "$QM"]
then
 echo "ERROR! No queue manager name supplied"
 exit 1
fi

if [-z "$TIMEOUT"]
then
 echo "ERROR! No timeout specified"
 exit 1
fi

for severity in immediate brutal
do
 # End the queue manager in the background to avoid
 # it blocking indefinitely. Run the TIMEOUT timer
 # at the same time to interrupt the attempt, and try a
 # more forceful version. If the brutal version fails,
 # nothing more can be done here.

 echo "Attempting ${severity} end of queue manager '${QM}'"
 case $severity in

 immediate)
 # Minimum severity of endmqm is immediate which severs connections.
 # HA cluster should not be delayed by clients
 endmqm -i ${QM} &
 ;;

 brutal)
 # This is a forced means of stopping queue manager processes.

 srchstr="(|-m)$QM *.*$"
 for process in amqzmuc0 amqzma0 amqfcxba amqfcpub amqpcsea amqzlaa0 \
 amqzlsa0 runmqchi runmqlsr amqcrsta amqrrmfa amqrmppa \
 amqzfuma amqzmuf0 amqzmur0 amqzmgr0
 do
 ps -ef | tr "\t" " " | grep $process | grep -v grep | \
 egrep "$srchstr" | awk '{print $2}' | \
 xargs kill -9 > /dev/null 2>&1
 done
 done
 done
```

```

esac

TIMED_OUT=yes
SECONDS=0
while (($SECONDS < ${TIMEOUT}))
do
 TIMED_OUT=yes
 i=0
 while [$i -lt 5]
 do
 # Check for execution controller termination
 srchstr="(|-m)$QM *.*$"
 cnt=`ps -ef | tr "\t" " " | grep amqzxa0 | grep -v grep | \
 egrep "$srchstr" | awk '{print $2}' | wc -l`
 i=`expr $i + 1`
 sleep 1
 if [$cnt -eq 0]
 then
 TIMED_OUT=no
 break
 fi
 done

 if [${TIMED_OUT} = "no"]
 then
 break
 fi

 echo "Waiting for ${severity} end of queue manager '${QM}'"
 sleep 1
done # timeout loop

if [${TIMED_OUT} = "yes"]
then
 continue # to next level of urgency
else
 break # queue manager is ended, job is done
fi

done # next phase

```

### Monitoring a queue manager:

It is usual to provide a way for the high availability (HA) cluster to monitor the state of the queue manager periodically. In most cases, you can use a shell script for this. Examples of suitable shell scripts are given here. You can tailor these scripts to your needs and use them to make additional monitoring checks specific to your environment.

From IBM WebSphere MQ Version 7.1, it is possible to have multiple installations of IBM MQ coexisting on a system. For more information about multiple installations, see [Multiple installations](#). If you intend to use the monitoring script across multiple installations, including installations at Version 7.1, or higher, you might need to perform some additional steps. If you have a primary installation, or you are using the script with versions earlier than Version 7.1, you do not need to specify the `MQ_INSTALLATION_PATH` to use the script. Otherwise, the following steps ensure that the `MQ_INSTALLATION_PATH` is identified correctly:

1. Use the `crtmqenv` command from a Version 7.1 installation to identify the correct `MQ_INSTALLATION_PATH` for a queue manager:

```
crtmqenv -m qmname
```

This command returns the correct `MQ_INSTALLATION_PATH` value for the queue manager specified by `qmname`.

2. Run the monitoring script with the appropriate `qmname` and `MQ_INSTALLATION_PATH` parameters.

**Note:** PowerHA for AIX does not provide a way of supplying a parameter to the monitoring program for the queue manager. You must create a separate monitoring program for each queue manager, that encapsulates the queue manager name. Here is an example of a script used on AIX to encapsulate the queue manager name:

```
#!/bin/ksh
su mqm -c name_of_monitoring_script qmname MQ_INSTALLATION_PATH
```

where *MQ\_INSTALLATION\_PATH* is an optional parameter that specifies the path to the installation of IBM MQ that the queue manager *qmname* is associated with.

The following script is not robust to the possibility that **runmqsc** hangs. Typically, HA clusters treat a hanging monitoring script as a failure and are themselves robust to this possibility.

The script does, however, tolerate the queue manager being in the starting state. This is because it is common for the HA cluster to start monitoring the queue manager as soon as it has started it. Some HA clusters distinguish between a starting phase and a running phase for resources, but it is necessary to configure the duration of the starting phase. Because the time taken to start a queue manager depends on the amount of work that it has to do, it is hard to choose a maximum time that starting a queue manager takes. If you choose a value that is too low, the HA cluster incorrectly assumes that the queue manager failed when it has not completed starting. This could result in an endless sequence of failovers.

This script must be run by the mqm user; it might therefore be necessary to wrap this script in a shell script to switch the user from the HA cluster user to mqm (an example shell script is provided in “Starting a queue manager under control of a high availability (HA) cluster” on page 1053 ):

```
#!/bin/ksh
#
This script tests the operation of the queue manager.
#
An exit code is generated by the runmqsc command:
0 => Either the queue manager is starting or the queue manager is running and responds.
Either is OK.
>0 => The queue manager is not responding and not starting.
#
This script must be run by the mqm user.
QM=$1
MQ_INSTALLATION_PATH=$2

if [-z "$QM"]
then
 echo "ERROR! No queue manager name supplied"
 exit 1
fi

if [-z "$MQ_INSTALLATION_PATH"]
then
 # No path specified, assume system primary install or MQ level < 7.1.0.0
 echo "INFO: Using shell default value for MQ_INSTALLATION_PATH"
else
 echo "INFO: Prefixing shell PATH variable with $MQ_INSTALLATION_PATH/bin"
 PATH=$MQ_INSTALLATION_PATH/bin:$PATH
fi

Test the operation of the queue manager. Result is 0 on success, non-zero on error.
echo "ping qmgr" | runmqsc ${QM} > /dev/null 2>&1
pingresult=$?

if [$pingresult -eq 0]
then # ping succeeded

 echo "Queue manager '${QM}' is responsive"
 result=0
```



```

else # ping failed

Don't condemn the queue manager immediately, it might be starting.
srchstr="(|-m)$QM *.*$"
cnt=`ps -ef | tr "\t" " " | grep strmqm | grep "$srchstr" | grep -v grep \
| awk '{print $2}' | wc -l`
if [$cnt -gt 0]
then
It appears that the queue manager is still starting up, tolerate
echo "Queue manager '${QM}' is starting"
result=0
else
There is no sign of the queue manager starting
echo "Queue manager '${QM}' is not responsive"
result=$pingresult
fi

fi

exit $result

```

### Putting the queue manager under control of the high availability (HA) cluster:

You must configure the queue manager, under control of the HA cluster, with the queue manager's IP address and shared disks.

To define a resource group to contain the queue manager and all of its associated resources, complete the following steps:

1. Create the resource group containing the queue manager, the queue manager's volume or disk group, and the queue manager's IP address. The IP address is a virtual IP address, not the IP address of the computer.
2. Verify that the HA cluster correctly switches the resources between the cluster nodes and is ready to control the queue manager.

### Deleting a queue manager from a high availability (HA) cluster node:

You might want to remove a queue manager from a node that is no longer required to run the queue manager.

To remove the queue manager from a node in an HA cluster, complete the following steps:

1. Remove the node from the HA cluster so that the HA cluster will no longer attempt to activate the queue manager on this node.
2. Use the following **rmvmqinf** command to remove the queue manager's configuration information:

```
rmvmqinf qmgrname
```

To completely delete the queue manager, use the **dltmqm** command. However, be aware that this completely deletes the queue manager's data and log files. When you have deleted the queue manager, you can use the **rmvmqinf** command to remove remaining configuration information from the other nodes.

## Supporting the Microsoft Cluster Service (MSCS)

Introducing and setting up MSCS to support failover of virtual servers.

This information applies to IBM MQ for Windows only.

The Microsoft Cluster Service (MSCS) enables you to connect servers into a *cluster*, giving higher availability of data and applications, and making it easier to manage the system. MSCS can automatically detect and recover from server or application failures.

MSCS supports *failover of virtual servers*, which correspond to applications, Web sites, print queues, or file shares (including, for example, their disk spindles, files, and IP addresses).

*Failover* is the process by which MSCS detects a failure in an application on one computer in the cluster, and shuts down the disrupted application in an orderly manner, transfers its state data to the other computer, and reinitiates the application there.

This section introduces MSCS clusters and describes setting up MSCS support in the following sections:

- “Introducing MSCS clusters”
- “Setting up IBM MQ for MSCS clustering” on page 1060

Then tells you how to configure IBM MQ for MSCS clustering, in the following sections:

- “Creating a queue manager for use with MSCS” on page 1062
- “Moving a queue manager to MSCS storage” on page 1063
- “Putting a queue manager under MSCS control” on page 1064
- “Removing a queue manager from MSCS control” on page 1070

And then gives some useful hints on using MSCS with IBM MQ, and details the IBM MQ MSCS support utility programs, in the following sections:

- “Hints and tips on using MSCS” on page 1071
- “IBM MQ MSCS support utility programs” on page 1074

### Introducing MSCS clusters:

MSCS clusters are groups of two or more computers, connected together and configured in such a way that, if one fails, MSCS performs a *failover*, transferring the state data of applications from the failing computer to another computer in the cluster and re-initiating their operation there.

“Using IBM MQ with high availability configurations” on page 1047 contains a comparison between MSCS clusters, multi-instance queue managers, and IBM MQ clusters.

In this section and its subordinate topics, the term *cluster*, when used by itself, **always** means an MSCS cluster. This is distinct from an IBM MQ cluster described elsewhere in this guide.

A two-machine cluster comprises two computers (for example, A and B) that are jointly connected to a network for client access using a *virtual IP address*. They might also be connected to each other by one or more private networks. A and B share at least one disk for the server applications on each to use. There is also another shared disk, which must be a redundant array of independent disks ( *RAID* ) Level 1, for the exclusive use of MSCS; this is known as the *quorum* disk. MSCS monitors both computers to check that the hardware and software are running correctly.

In a simple setup such as this, both computers have all the applications installed on them, but only computer A runs with live applications; computer B is just running and waiting. If computer A encounters any one of a range of problems, MSCS shuts down the disrupted application in an orderly

manner, transfers its state data to the other computer, and re-initiates the application there. This is known as a *failover*. Applications can be made *cluster-aware* so that they interact fully with MSCS and failover gracefully.

A typical setup for a two-computer cluster is as shown in Figure 153.

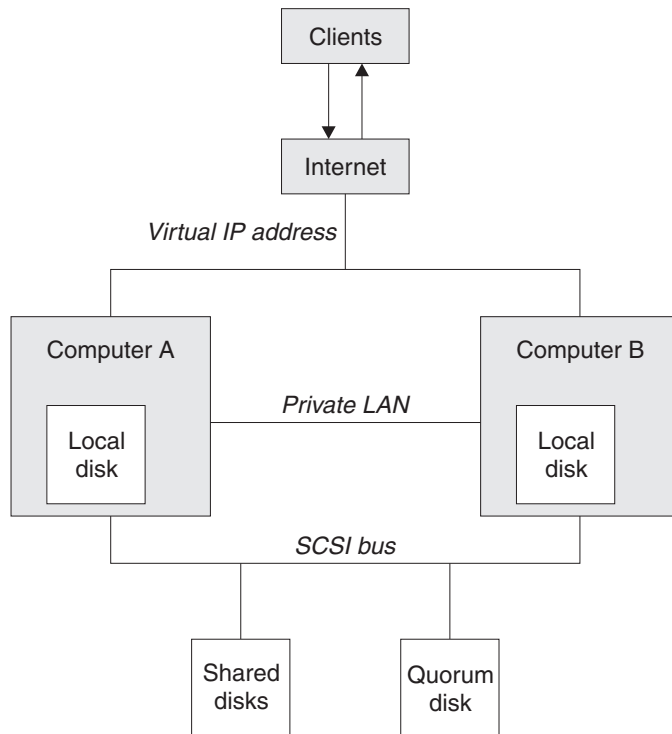


Figure 153. Two-computer MSCS cluster

Each computer can access the shared disk, but only one at a time, under the control of MSCS. In the event of a failover, MSCS switches the access to the other computer. The shared disk itself is usually a RAID, but need not be.

Each computer is connected to the external network for client access, and each has an IP address. However an external client, communicating with this cluster, is aware of only one *virtual IP address*, and MSCS routes the IP traffic within the cluster appropriately.

MSCS also performs its own communications between the two computers, either over one or more private connections or over the public network, for example to monitor their states using the heartbeat, and to synchronize their databases.

## Setting up IBM MQ for MSCS clustering:

You configure IBM MQ for clustering by making the queue manager the unit of failover to MSCS. You define a queue manager as a resource to MSCS, which can then monitor it, and transfer it to another computer in the cluster if there is a problem.

To set your system up for this, you start by installing IBM MQ on each computer in the cluster.

As the queue manager is associated with the IBM MQ installation name, the IBM MQ installation name on all the computers in the cluster should be the same. See *Installing and uninstalling*.

The queue managers themselves need to exist only on the computer on which you create them. In the event of a failover, the MSCS initiates the queue managers on the other computer. The queue managers, however, must have their log and data files on a cluster shared disk, and not on a local drive. If you have a queue manager already installed on a local drive, you can migrate it using a tool provided with IBM MQ; see *“Moving a queue manager to MSCS storage”* on page 1063. If you want to create new queue managers for use with MSCS, see *“Creating a queue manager for use with MSCS”* on page 1062.

After installation and migration, use the MSCS Cluster Administrator to make MSCS aware of your queue managers; see *“Putting a queue manager under MSCS control”* on page 1064.

If you decide to remove a queue manager from MSCS control, use the procedure described in *“Removing a queue manager from MSCS control”* on page 1070.

### *Setup symmetry:*

When an application switches from one node to the other it must behave in the same way, regardless of node. The best way of ensuring this is to make the environments identical.

If you can, set up a cluster with identical hardware, operating system software, product software, and configuration on each computer. In particular, ensure that all the required software installed on the two computers is identical in terms of version, maintenance level, SupportPacs, paths and exits, and that there is a common namespace (security environment) as described in *“MSCS security.”*

### *MSCS security:*

For successful MSCS security, follow these guidelines.

The guidelines are as follows:

- Make sure you that you have identical software installations on each computer in the cluster.
- Create a common namespace (security environment) across the cluster.
- Make the nodes of the MSCS cluster members of a domain, within which the user account that is the *cluster owner* is a domain account.
- Make the other user accounts on the cluster also domain accounts, so that they are available on both nodes. This is automatically the case if you already have a domain, and the accounts relevant to IBM MQ are domain accounts. If you do not currently have a domain, consider setting up a *mini-domain* to cater for the cluster nodes and relevant accounts. Your aim is to make your cluster of two computers look like a single computing resource.

Remember that an account that is local to one computer does not exist on the other one. Even if you create an account with the same name on the other computer, its security identifier (SID) is different, so, when your application is moved to the other node, the permissions do not exist on that node.

During a failover or move, IBM MQ MSCS support ensures that all files that contain queue manager objects have equivalent permissions on the destination node. Explicitly, the code checks that the

Administrators and mqm groups, and the SYSTEM account, have full control, and that if Everyone had read access on the old node, that permission is added on the destination node.

You can use a domain account to run your IBM MQ Service. Make sure that it exists in the local mqm group on each computer in the cluster.

*Using multiple queue managers with MSCS:*

If you are running more than one queue manager on a computer, you can choose one of these setups.

The setups are as follows:

- All the queue managers in a single group. In this configuration, if a problem occurs with any queue manager, all the queue managers in the group failover to the other computer as a group.
- A single queue manager in each group. In this configuration, if a problem occurs with the queue manager, it alone fails over to the other computer without affecting the other queue managers.
- A mixture of the first two setups.

*Cluster modes:*

There are two modes in which you might run a cluster system with IBM MQ: Active/Passive or Active/Active.

**Note:** If you are using MSCS together with the Microsoft Transaction Server (COM+), you cannot use Active/Active mode.

### **Active/Passive mode**

In Active/Passive mode, computer A has the running application on it, and computer B is backup, only being used when MSCS detects a problem.

You can use this mode with only one shared disk, but, if any application causes a failover, **all** the applications must be transferred as a group (because only one computer can access the shared disk at a time).

You can configure MSCS with A as the *preferred* computer. Then, when computer A has been repaired or replaced and is working properly again, MSCS detects this and automatically switches the application back to computer A.

If you run more than one queue manager, consider having a separate shared disk for each. Then put each queue manager in a separate group in MSCS. In this way, any queue manager can failover to the other computer without affecting the other queue managers.

### **Active/Active mode**

In Active/Active mode, computers A and B both have running applications, and the groups on each computer are set to use the other computer as backup. If a failure is detected on computer A, MSCS transfers the state data to computer B, and reinitiates the application there. computer B then runs its own application and A's.

For this setup you need at least two shared disks. You can configure MSCS with A as the preferred computer for A's applications, and B as the preferred computer for B's applications. After failover and repair, each application automatically ends up back on its own computer.

For IBM MQ this means that you could, for example, run two queue managers, one on each of A and B, with each exploiting the full power of its own computer. After a failure on computer A, both queue

managers will run on computer B. This will mean sharing the power of the one computer, with a reduced ability to process large quantities of data at speed. However, your critical applications will still be available while you find and repair the fault on A.

### **Creating a queue manager for use with MSCS:**

This procedure ensures that a new queue manager is created in such a way that it is suitable for preparing and placing under MSCS control.

You start by creating the queue manager with all its resources on a local drive, and then migrate the log files and data files to a shared disk. (You can reverse this operation.) Do **not** attempt to create a queue manager with its resources on a shared drive.

You can create a queue manager for use with MSCS in two ways, either from a command prompt, or in the IBM MQ Explorer. The advantage of using a command prompt is that the queue manager is created *stopped* and set to *manual startup*, which is ready for MSCS. (The IBM MQ Explorer automatically starts a new queue manager and sets it to automatic startup after creation. You have to change this.)

### **Creating a queue manager from a command prompt**

Follow these steps to create a queue manager from a command prompt, for use with MSCS:

1. Ensure that you have the environment variable MQSPREFIX set to refer to a local drive, for example C:\WebSphere MQ. If you change this, reboot the machine so that the System account picks up the change. If you do not set the variable, the queue manager is created in the IBM MQ default directory for queue managers.
2. Create the queue manager using the **crtmqm** command. For example, to create a queue manager called `mscs_test` in the default directory, use:  
`crtmqm mscs_test`
3. Proceed to “Moving a queue manager to MSCS storage” on page 1063.

### **Creating a queue manager using the IBM MQ Explorer**

Follow these steps to create a queue manager using the IBM MQ Explorer, for use with MSCS:

1. Start the IBM MQ Explorer from the Start menu.
2. In the Navigator View, expand the tree nodes to find the **Queue Managers** tree node.
3. Right-click the **Queue Managers** tree node, and select **New -> Queue Manager** . The Create Queue Manager panel is displayed.
4. Complete the dialog (Step 1), then click **Next>** .
5. Complete the dialog (Step 2), then click **Next>** .
6. Complete the dialog (Step 3), ensuring that **Start Queue Manager** and **Create Server Connection Channel** are not selected, then click **Next>** .
7. Complete the dialog (Step 4), then click **Finish** .
8. Proceed to “Moving a queue manager to MSCS storage” on page 1063.

## Moving a queue manager to MSCS storage:

This procedure configures an existing queue manager to make it suitable for putting under MSCS control.

To achieve this, you move the log files and data files to shared disks to make them available to the other computer in the event of a failure. For example, the existing queue manager might have paths such as C:\WebSphere MQ\log\<QMname> and C:\WebSphere MQ\qmgrs\<QMname>. Do *not* try to move the files by hand; use the utility program supplied as part of IBM MQ MSCS Support as described in this topic.

If the queue manager being moved uses SSL connections and the SSL key repository is in the queue manager data directory on the local machine, then the key repository will be moved with the rest of the queue manager to the shared disk. By default, the queue manager attribute that specifies the SSL key repository location, `SSLKEYR`, is set to `MQ_INSTALLATION_PATH\qmgrs\QMGRNAME\ssl\key`, which is under the queue manager data directory. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed. The `hamvmqm` command does not modify this queue manager attribute. In this situation you must modify the queue manager attribute, `SSLKEYR`, using the IBM MQ Explorer or the MQSC command `ALTER QMGR`, to point to the new SSL key repository file.

The procedure is as follows:

1. Shut down the queue manager, and check that there are no errors.
2. If the queue manager's log files or queue files are already stored on a shared disk, skip the rest of this procedure and proceed directly to "Putting a queue manager under MSCS control" on page 1064.
3. Make a full media backup of the queue files and log files and store the backup in a safe place (see "Queue manager log files" on page 1072 for why this is important).
4. If you already have a suitable shared disk resource proceed to step 6. Otherwise, using the MSCS Cluster Administrator to create a resource of type *shared disk* with sufficient capacity to store the queue manager log files and data (queue) files.
5. Test the shared disk by using the MSCS Cluster Administrator to move it from one cluster node to the other and back again.
6. Make sure that the shared disk is online on the cluster node where the queue manager log and data files are stored locally.
7. Run the utility program to move the queue manager as follows:

```
hamvmqm /m qmname /dd " e: \ IBM MQ " /ld " e: \ IBM MQ \log"
```

substituting your queue manager name for *qmname*, your shared disk drive letter for *e*, and your chosen directory for *IBM MQ*. The directories are created if they do not already exist.

8. Test the queue manager to ensure that it works, using the IBM MQ Explorer. For example:
  - a. Right-click the queue manager tree node, then select **Start** . The queue manager starts.
  - b. Right-click the **Queues** tree node, then select **New -> Local Queue...** , and give the queue a name.
  - c. Click **Finish** .
  - d. Right-click the queue, then select **Put Test Message...** . The Put Test Message panel is displayed.
  - e. Type some message text, then click **Put Test Message** , and close the panel.
  - f. Right-click the queue, then select **Browse Messages...** . The Message Browser panel is displayed.
  - g. Ensure your message is on the queue, then click **Close** . The Message Browser panel closes.
  - h. Right-click the queue, then select **Clear Messages...** . The messages on the queue are cleared.
  - i. Right-click the queue, then select **Delete...** . A confirmation panel is displayed, click **OK** . The queue is deleted.
  - j. Right-click the queue manager tree node, then select **Stop...** . The End Queue Manager panel is displayed.

- k. Click **OK** . The queue manager stops.
9. As IBM MQ Administrator ensure that the startup attribute of the queue manager is set to manual. In the IBM MQ Explorer, set the Startup field to manual in the queue manager properties panel.
10. Proceed to "Putting a queue manager under MSCS control."

### Putting a queue manager under MSCS control:

The tasks involved in placing a queue manager under MSCS control, including prerequisite tasks.

### Before you put a queue manager under MSCS control

Before you put a queue manager under MSCS control, perform the following tasks:

1. Ensure that IBM MQ and its MSCS Support are installed on both machines in the cluster and that the software on each computer is identical, as described in "Setting up IBM MQ for MSCS clustering" on page 1060.
2. Use the **haregtyp** utility program to register IBM MQ as an MSCS resource type on all the cluster nodes. See "IBM MQ MSCS support utility programs" on page 1074 for additional information.
3. If you have not yet created the queue manager, see "Creating a queue manager for use with MSCS" on page 1062.
4. If you have created the queue manager, or it already exists, ensure that you have carried out the procedure in "Moving a queue manager to MSCS storage" on page 1063.
5. Stop the queue manager, if it is running, using either a command prompt or the IBM MQ Explorer.
6. Test MSCS operation of the shared drives before going on to either of the following Windows procedures in this topic.

### Windows Server 2012

To place a queue manager under MSCS control on Windows Server 2012, use the following procedure:

1. Log in to the cluster node computer hosting the queue manager, or log in to a remote workstation as a user with cluster administration permissions, and connect to the cluster node hosting the queue manager.
2. Start the Failover Cluster Management tool.
3. Right-click **Failover Cluster Management > Connect Cluster ...** to open a connection to the cluster.
4. In contrast to the group scheme used in the MSCS Cluster Administrator on previous versions of Windows, the Failover Cluster Management tool uses the concept of services and applications. A configured service or application contains all the resources necessary for one application to be clustered. You can configure a queue manager under MSCS as follows:
  - a. Right-click on the cluster and select **Configure Role** to start the configuration wizard.
  - b. Select **Other Server** on the "Select Service or Application" panel.
  - c. Select an appropriate IP address as a client access point.

This address should be an unused IP address to be used by clients and other queue managers to connect to the *virtual* queue manager. This IP address is not the normal (static) address of either node; it is an additional address that *floats* between them. Although MSCS handles the routing of this address, it does **not** verify that the address can be reached.
  - d. Assign a storage device for exclusive use by the queue manager. This device needs to be created as a resource instance before it can be assigned.

You can use one drive to store both the logs and queue files, or you can split them up across drives. In either case, if each queue manager has its own shared disk, ensure that all drives used by this queue manager are exclusive to this queue manager, that is, that nothing else relies on the drives. Also ensure that you create a resource instance for every drive that the queue manager uses.



The resource type for a drive depends on the SCSI support you are using; refer to your SCSI adapter instructions. There might already be groups and resources for each of the shared drives. If so, you do not need to create the resource instance for each drive. Move it from its current group to the one created for the queue manager.

For each drive resource, set possible owners to both nodes. Set dependent resources to none.

- e. Select the **MQSeries MSCS** resource on the "Select Resource Type" panel.
  - f. Complete the remaining steps in the wizard.
5. Before bringing the resource online, the MQSeries MSCS resource needs additional configuration:
- a. Select the newly defined service which contains a resource called 'New MQSeries MSCS'.
  - b. Right-click **Properties** on the MQ resource.
  - c. Configure the resource:
    - Name ; choose a name that makes it easy to identify which queue manager it is for.
    - Run in a separate Resource Monitor ; for better isolation
    - Possible owners ; set both nodes
    - Dependencies ; add the drive and IP address for this queue manager.

**Warning:** Failure to add these dependencies means that IBM MQ attempts to write the queue manager status to the wrong cluster disk during failovers. Because many processes might be attempting to write to this disk simultaneously, some IBM MQ processes could be blocked from running.

- Parameters ; as follows:
  - QueueManagerName (required); the name of the queue manager that this resource is to control. This queue manager must exist on the local computer.
  - PostOnlineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from offline to online. For more details see "PostOnlineCommand and PreOfflineCommand" on page 1073.
  - PreOfflineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from online to offline. For more details see "PostOnlineCommand and PreOfflineCommand" on page 1073.

**Note:** The *looksAlive* poll interval is set to default value of 5000 ms. The *isAlive* poll interval is set to default value of 60000 ms. These defaults can only be modified after the resource definition has been completed. For further details see "Summary of looksAlive and isAlive polling" on page 1069.

- d. Optionally, set a preferred node (but note the comments in "Using preferred nodes" on page 1074 )
  - e. The *Failover Policy* is set by default to sensible values, but you can tune the thresholds and periods that control *Resource Failover* and *Group Failover* to match the loads placed on the queue manager.
6. Test the queue manager by bringing it online in the MSCS Cluster Administrator and subjecting it to a test workload. If you are experimenting with a test queue manager, use the IBM MQ Explorer. For example:
- a. Right-click the **Queues** tree node, then select **New -> Local Queue...** , and give the queue a name.
  - b. Click **Finish** . The queue is created, and displayed in the content view.
  - c. Right-click the queue, then select **Put Test Message...** . The Put Test Message panel is displayed.
  - d. Type some message text, then click **Put Test Message** , and close the panel.
  - e. Right-click the queue, then select **Browse Messages...** . The Message Browser panel is displayed.
  - f. Ensure that your message is on the queue, then click **Close** . The Message Browser panel closes.
  - g. Right-click the queue, then select **Clear Messages...** . The messages on the queue are cleared.

- h. Right-click the queue, then select **Delete...** . A confirmation panel is displayed, click **OK** . The queue is deleted.
7. Test that the queue manager can be taken offline and back online using the MSCS Cluster Administrator.
8. Simulate a failover.
 

In the MSCS Cluster Administrator, right-click the group containing the queue manager and select **Move Group**. This can take some minutes to do. (If at other times you want to move a queue manager to another node quickly, follow the procedure in “Moving a queue manager to MSCS storage” on page 1063.) You can also right-click and select **Initiate Failure** ; the action (local restart or failover) depends on the current state and the configuration settings.

## Windows Server 2008

To place a queue manager under MSCS control on Windows Server 2008, use the following procedure:

1. Log in to the cluster node computer hosting the queue manager, or log in to a remote workstation as a user with cluster administration permissions, and connect to the cluster node hosting the queue manager.
2. Start the Failover Cluster Management tool.
3. Right-click **Failover Cluster Management > Manage a Cluster ...** to open a connection to the cluster.
4. In contrast to the group scheme used in the MSCS Cluster Administrator on previous versions of Windows, the Failover Cluster Management tool uses the concept of services and applications. A configured service or application contains all the resources necessary for one application to be clustered. You can configure a queue manager under MSCS as follows:
  - a. Right-click **Services and Applications > Configure a Service or Application ...** to start the configuration wizard.
  - b. Select **Other Server** on the Select Service or Application panel.
  - c. Select an appropriate IP address as a client access point.
 

This address should be an unused IP address to be used by clients and other queue managers to connect to the *virtual* queue manager. This IP address is not the normal (static) address of either node; it is an additional address that *floats* between them. Although MSCS handles the routing of this address, it does **not** verify that the address can be reached.
  - d. Assign a storage device for exclusive use by the queue manager. This device needs to be created as a resource instance before it can be assigned.
 

You can use one drive to store both the logs and queue files, or you can split them up across drives. In either case, if each queue manager has its own shared disk, ensure that all drives used by this queue manager are exclusive to this queue manager, that is, that nothing else relies on the drives. Also ensure that you create a resource instance for every drive that the queue manager uses.

The resource type for a drive depends on the SCSI support you are using; refer to your SCSI adapter instructions. There might already be groups and resources for each of the shared drives. If so, you do not need to create the resource instance for each drive. Move it from its current group to the one created for the queue manager.

For each drive resource, set possible owners to both nodes. Set dependent resources to none.
  - e. Select the **MQSeries MSCS** resource on the Select Resource Type panel.
  - f. Complete the remaining steps in the wizard.
5. Before bringing the resource online, the MQSeries MSCS resource needs additional configuration:
  - a. Select the newly defined service which contains a resource called 'New MQSeries MSCS'.
  - b. Right-click **Properties** on the MQ resource.
  - c. Configure the resource:
    - Name ; choose a name that makes it easy to identify which queue manager it is for.

- Run in a separate Resource Monitor ; for better isolation
- Possible owners ; set both nodes
- Dependencies ; add the drive and IP address for this queue manager.

**Warning:** Failure to add these dependencies means that IBM MQ attempts to write the queue manager status to the wrong cluster disk during failovers. Because many processes might be attempting to write to this disk simultaneously, some IBM MQ processes could be blocked from running.

- Parameters ; as follows:
  - QueueManagerName (required); the name of the queue manager that this resource is to control. This queue manager must exist on the local computer.
  - PostOnlineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from offline to online. For more details see “PostOnlineCommand and PreOfflineCommand” on page 1073.
  - PreOfflineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from online to offline. For more details see “PostOnlineCommand and PreOfflineCommand” on page 1073.

**Note:** The *looksAlive* poll interval is set to default value of 5000 ms. The *isAlive* poll interval is set to default value of 60000 ms. These defaults can only be modified after the resource definition has been completed. For further details see “Summary of looksAlive and isAlive polling” on page 1069.

- d. Optionally, set a preferred node (but note the comments in “Using preferred nodes” on page 1074 )
  - e. The *Failover Policy* is set by default to sensible values, but you can tune the thresholds and periods that control *Resource Failover* and *Group Failover* to match the loads placed on the queue manager.
6. Test the queue manager by bringing it online in the MSCS Cluster Administrator and subjecting it to a test workload. If you are experimenting with a test queue manager, use the IBM MQ Explorer. For example:
    - a. Right-click the **Queues** tree node, then select **New -> Local Queue...** , and give the queue a name.
    - b. Click **Finish** . The queue is created, and displayed in the content view.
    - c. Right-click the queue, then select **Put Test Message...** . The Put Test Message panel is displayed.
    - d. Type some message text, then click **Put Test Message** , and close the panel.
    - e. Right-click the queue, then select **Browse Messages...** . The Message Browser panel is displayed.
    - f. Ensure that your message is on the queue, then click **Close** . The Message Browser panel closes.
    - g. Right-click the queue, then select **Clear Messages...** . The messages on the queue are cleared.
    - h. Right-click the queue, then select **Delete...** . A confirmation panel is displayed, click **OK** . The queue is deleted.
  7. Test that the queue manager can be taken offline and back online using the MSCS Cluster Administrator.
  8. Simulate a failover.
 

In the MSCS Cluster Administrator, right-click the group containing the queue manager and select **Move Group**. This can take some minutes to do. (If at other times you want to move a queue manager to another node quickly, follow the procedure in “Moving a queue manager to MSCS storage” on page 1063.) You can also right-click and select **Initiate Failure** ; the action (local restart or failover) depends on the current state and the configuration settings.

## Windows 2003

To place a queue manager under MSCS control on Windows 2003, use the following procedure:

1. Log in to the cluster node computer hosting the queue manager, or log in to a remote workstation as a user with cluster administration permissions, and connect to the cluster node hosting the queue manager.
2. Start the MSCS Cluster Administrator.
3. Open a connection to the cluster.
4. Create an MSCS group to be used to contain the resources for the queue manager. Name the group in such a way that it is obvious which queue manager it relates to. Each group can contain multiple queue managers, as described in “Using multiple queue managers with MSCS” on page 1061.

Use the group for all the remaining steps.

5. Create a resource instance for each of the SCSI logical drives that the queue manager uses. You can use one drive to store both the logs and queue files, or you can split them up across drives. In either case, if each queue manager has its own shared disk, ensure that all drives used by this queue manager are exclusive to this queue manager, that is, that nothing else relies on the drives. Also ensure that you create a resource instance for every drive that the queue manager uses. The resource type for a drive depends on the SCSI support you are using; refer to your SCSI adapter instructions. There might already be groups and resources for each of the shared drives. If so, you do not need to create the resource instance for each drive. Move it from its current group to the one created for the queue manager.

For each drive resource, set possible owners to both nodes. Set dependent resources to none.

6. Create a resource instance for the IP address. Create an IP address resource (resource type *IP address* ). This address should be an unused IP address to be used by clients and other queue managers to connect to the *virtual* queue manager. This IP address is not the normal (static) address of either node; it is an additional address that *floats* between them. Although MSCS handles the routing of this address, it does **not** verify that the address can be reached.
7. Create a resource instance for the queue manager.

Create a resource of type *IBM MQ MSCS*. The wizard prompts you for various items, including the following:

- Name ; choose a name that makes it easy to identify which queue manager it is for.
- Add to group ; use the group that you created
- Run in a separate Resource Monitor ; for better isolation
- Possible owners ; set both nodes
- Dependencies ; add the drive and IP address for this queue manager.

**Warning:** Failure to add these dependencies means that IBM MQ attempts to write the queue manager status to the wrong cluster disk during failovers. Because many processes might be attempting to write to this disk simultaneously, some IBM MQ processes could be blocked from running.

- Parameters ; as follows:
  - QueueManagerName (required); the name of the queue manager that this resource is to control. This queue manager must exist on the local computer.
  - PostOnlineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from offline to online. For more details see “PostOnlineCommand and PreOfflineCommand” on page 1073.
  - PreOfflineCommand (optional); you can specify a program to run whenever the queue manager resource changes its state from online to offline. For more details see “PostOnlineCommand and PreOfflineCommand” on page 1073.

**Note:** The *looksAlive* poll interval is set to default value of 5000 ms. The *isAlive* poll interval is set to default value of 30000 ms. These defaults can only be modified after the resource definition has been completed. For further details see “Summary of looksAlive and isAlive polling.”

8. Optionally, set a preferred node (but note the comments in “Using preferred nodes” on page 1074 )
9. The *Failover Policy* (as defined in the properties for the group) is set by default to sensible values, but you can tune the thresholds and periods that control *Resource Failover* and *Group Failover* to match the loads placed on the queue manager.
10. Test the queue manager by bringing it online in the MSCS Cluster Administrator and subjecting it to a test workload. If you are experimenting with a test queue manager, use the IBM MQ Explorer. For example:
  - a. Right-click the **Queues** tree node, then select **New -> Local Queue...** , and give the queue a name.
  - b. Click **Finish** . The queue is created, and displayed in the content view.
  - c. Right-click the queue, then select **Put Test Message...** . The Put Test Message panel is displayed.
  - d. Type some message text, then click **Put Test Message** , and close the panel.
  - e. Right-click the queue, then select **Browse Messages...** . The Message Browser panel is displayed.
  - f. Ensure that your message is on the queue, then click **Close** . The Message Browser panel closes.
  - g. Right-click the queue, then select **Clear Messages...** . The messages on the queue are cleared.
  - h. Right-click the queue, then select **Delete...** . A confirmation panel is displayed, click **OK** . The queue is deleted.
11. Test that the queue manager can be taken offline and back online using the MSCS Cluster Administrator.
12. Simulate a failover.

In the MSCS Cluster Administrator, right-click the group containing the queue manager and select Move Group. This can take some minutes to do. (If at other times you want to move a queue manager to another node quickly, follow the procedure in “Moving a queue manager to MSCS storage” on page 1063.) You can also right-click and select Initiate Failure ; the action (local restart or failover) depends on the current state and the configuration settings.

### Summary of looksAlive and isAlive polling:

*looksAlive* and *isAlive* are intervals at which MSCS calls back into the resource types supplied library code and requests that the resource performs checks to determine the working status of itself. This ultimately determines if MSCS attempts to fail over the resource.

On every occasion that the *looksAlive* interval elapses (default 5000 ms), the queue manager resource is called to perform its own check to determine if its status is satisfactory.

On every occasion that the *isAlive* interval elapses (default 30000 ms), another call is made to the queue manager resource for it to perform another check to determine if the resource is functioning correctly. This enables two levels of resource type checking.

1. A *looksAlive* status check to establish if the resource appears to be functioning.
2. A more significant *isAlive* check that determines if the queue manager resource is active.

If the queue manager resource is determined not to be active, MSCS, based on other advanced MSCS options, triggers a fail over for the resource and associated dependant resources to another node in the cluster. For further information, see MSCS documentation.

## Removing a queue manager from MSCS control:

You can remove queue managers from MSCS control, and return them to manual administration.

You do not need to remove queue managers from MSCS control for maintenance operations. You can do that by taking a queue manager offline temporarily, using the MSCS Cluster Administrator. Removing a queue manager from MSCS control is a more permanent change; only do it if you decide that you no longer want MSCS to have any further control of the queue manager.

If the queue manager is being removed uses TSL or SSL connections you must modify the queue manager attribute, `SSLKEYR`, using the IBM MQ Explorer or the MQSC command `ALTER QMGR`, to point to the SSL key repository file on the local directory.

The procedure is:

1. Take the queue manager resource offline using the MSCS Cluster Administrator, as described in "Taking a queue manager offline from MSCS"
2. Destroy the resource instance. This does not destroy the queue manager.
3. Optionally, migrate the queue manager files back from shared drives to local drives. To do this, see "Returning a queue manager from MSCS storage."
4. Test the queue manager.

## Taking a queue manager offline from MSCS

To take a queue manager offline from MSCS, perform the following steps:

1. Start the MSCS Cluster Administrator.
2. Open a connection to the cluster.
3. Select Groups, or Role if you are using Windows 2012, and open the group containing the queue manager to be moved.
4. Select the queue manager resource.
5. Right-click it and select Offline.
6. Wait for completion.

## Returning a queue manager from MSCS storage

This procedure configures the queue manager to be back on its computer's local drive, that is, it becomes a *normal* IBM MQ queue manager. To achieve this, you move the log files and data files from the shared disks. For example, the existing queue manager might have paths such as `E:\WebSphere MQ\log\<QMname>` and `E:\WebSphere MQ\qmgrs\<QMname>`. Do not try to move the files by hand; use the **hamvmqm** utility program supplied as part of IBM MQ MSCS Support:

1. Make a full media backup of the queue files and log files and store the backup in a safe place (see "Queue manager log files" on page 1072 for why this is important).
2. Decide which local drive to use and ensure that it has sufficient capacity to store the queue manager log files and data (queue) files.
3. Make sure that the shared disk on which the files currently reside is online on the cluster node to which to move the queue manager log and data files.
4. Run the utility program to move the queue manager as follows:

```
hamvmqm /m qmname /dd " c:\ IBM MQ " /ld "c:\ IBM MQ \log"
```

substituting your queue manager name for *qmname*, your local disk drive letter for *c*, and your chosen directory for *IBM MQ* (the directories are created if they do not already exist).

5. Test the queue manager to ensure that it works (as described in "Moving a queue manager to MSCS storage" on page 1063 ).

## Hints and tips on using MSCS:

This section contains some general information to help you use IBM MQ support for MSCS effectively.

This section contains some general information to help you use IBM MQ support for MSCS effectively.

How long does it take to fail a queue manager over from one machine to the other? This depends heavily on the amount of workload on the queue manager and on the mix of traffic, for example, how much of it is persistent, within sync point, and how much committed before the failure. IBM tests have given failover and failback times of about a minute. This was on a very lightly loaded queue manager and actual times will vary considerably depending on load.

*Verifying that MSCS is working:*

Follow these steps to ensure that you have a running MSCS cluster.

The task descriptions starting with “Creating a queue manager for use with MSCS” on page 1062 assume that you have a running MSCS cluster within which you can create, migrate, and destroy resources. If you want to make sure that you have such a cluster:

1. Using the MSCS Cluster Administrator, create a group.
2. Within that group, create an instance of a generic application resource, specifying the system clock (path name C:\winnt\system32\clock.exe and working directory of C:\).
3. Make sure that you can bring the resource online, that you can move the group that contains it to the other node, and that you can take the resource offline.

*Manual startup:*

For a queue manager managed by MSCS, you *must* set the startup attribute to manual. This ensures that the IBM MQ MSCS support can restart the MQSeries Service without immediately starting the queue manager.

The IBM MQ MSCS support needs to be able to restart the service so that it can perform monitoring and control, but must itself remain in control of which queue managers are running, and on which machines. See “Moving a queue manager to MSCS storage” on page 1063 for more information.

*MSCS and queue managers:*

Considerations concerning queue managers when using MSCS.

### Creating a matching queue manager on the other node

For clustering to work with IBM MQ, you need an identical queue manager on node B for each one on node A. However, you do not need to explicitly create the second one. You can create or prepare a queue manager on one node, move it to the other node as described in “Moving a queue manager to MSCS storage” on page 1063, and it is fully duplicated on that node.

### Default queue managers

Do not use a default queue manager under MSCS control. A queue manager does not have a property that makes it the default; IBM MQ keeps its own separate record. If you move a queue manager set to be the default to the other computer on failover, it does not become the default there. Make all your applications refer to specific queue managers by name.

## Deleting a queue manager

Once a queue manager has moved node, its details exist in the registry on both computers. When you want to delete it, do so as normal on one computer, and then run the utility described in “IBM MQ MSCS support utility programs” on page 1074 to clean up the registry on the other computer.

## Support for existing queue managers

You can put an existing queue manager under MSCS control, provided that you can put your queue manager log files and queue files on a disk that is on the shared SCSI bus between the two machines (see Figure 153 on page 1059 ). You need to take the queue manager offline briefly while the MSCS Resource is created.

If you want to create a new queue manager, create it independently of MSCS, test it, then put it under MSCS control. See:

- “Creating a queue manager for use with MSCS” on page 1062
- “Moving a queue manager to MSCS storage” on page 1063
- “Putting a queue manager under MSCS control” on page 1064

## Telling MSCS which queue managers to manage

You choose which queue managers are placed under MSCS control by using the MSCS Cluster Administrator to create a resource instance for each such queue manager. This process presents you with a list of resources from which to select the queue manager that you want that instance to manage.

## Queue manager log files

When you move a queue manager to MSCS storage, you move its log and data files to a shared disk (for an example see “Moving a queue manager to MSCS storage” on page 1063 ).

It is advisable before you move, to shut the queue manager cleanly and take a full backup of the data files and log files.

## Multiple queue managers

IBM MQ MSCS support allows you to run multiple queue managers on each machine and to place individual queue managers under MSCS control.

*Always use MSCS to manage clusters:*

Do not try to perform start and stop operations directly on any queue manager under the control of MSCS, using either the control commands or the IBM MQ Explorer. Instead, use MSCS Cluster Administrator to bring the queue manager online or take it offline.

Using the MSCS Cluster Administrator is partly to prevent possible confusion caused by MSCS reporting that the queue manager is offline, when in fact you have started it outside the control of MSCS. More seriously, stopping a queue manager without using MSCS is detected by MSCS as a failure, initiating failover to the other node.



*Working in Active/Active mode:*

Both computers in the MSCS cluster can run queue managers in Active/Active mode. You do not need to have a completely idle machine acting as standby (but you can, if you want, in Active/Passive Mode).

If you plan to use both machines to run workload, provide each with sufficient capacity (processor, memory, secondary storage) to run the entire cluster workload at a satisfactory level of performance.

**Note:** If you are using MSCS together with Microsoft Transaction Server (COM+), you **cannot** use Active/Active mode. This is because, to use IBM MQ with MSCS and COM+:

- Application components that use IBM MQ COM+ support must run on the same computer as the Distributed Transaction Coordinator (DTC), a part of COM+.
- The queue manager must also run on the same computer.
- The DTC must be configured as an MSCS resource, and can therefore run on only one of the computers in the cluster at any time.

*PostOnlineCommand and PreOfflineCommand:*

Use these commands to integrate IBM MQ MSCS support with other systems. You can use them to issue IBM MQ commands, with some restrictions.

Specify these commands in the Parameters to a resource of type IBM MQ MSCS. You can use them to integrate IBM MQ MSCS support with other systems or procedures. For example, you could specify the name of a program that sends a mail message, activates a pager, or generates some other form of alert to be captured by another monitoring system.

PostOnlineCommand is invoked when the resource changes from offline to online; PreOfflineCommand is invoked for a change from online to offline. When invoked these commands are run, by default, from the Windows system directory. Because IBM MQ uses a 32-bit resource monitor process, on Windows 64-bit systems, this is the \Windows\SysWOW64 directory rather than the \Windows\system32 directory. For more information, see the Microsoft documentation about file redirection in a Windows x64 environment. Both commands run under the user account used to run the MSCS Cluster Service; and are invoked asynchronously; IBM MQ MSCS support does not wait for them to complete before continuing. This eliminates any risk that they might block or delay further cluster operations.

You can also use these commands to issue IBM MQ commands, for example to restart Requester channels. However, the commands are run at the point in time when the queue manager's state changes so they are not intended to perform long-running functions and must not make assumptions about the current state of the queue manager; it is quite possible that, immediately after the queue manager was brought online, an administrator issued an offline command.

If you want to run programs that depend on the state of the queue manager, consider creating instances of the MSCS Generic Application resource type, placing them in the same MSCS group as the queue manager resource, and making them dependent on the queue manager resource.

*Using preferred nodes:*

It can be useful when using Active/Active mode to configure a *preferred node* for each queue manager. However, in general it is better not to set a preferred node but to rely on a manual failback.

Unlike some other relatively stateless resources, a queue manager can take a while to fail over (or back) from one node to the other. To avoid unnecessary outages, test the recovered node before failing a queue manager back to it. This precludes use of the *immediate failback* setting. You can configure failback to occur between certain times of day.

Probably the safest route is to move the queue manager back manually to the required node, when you are certain that the node is fully recovered. This precludes use of the preferred node option.

*If COM+ errors occur in the Application Event log:*

When you install IBM MQ on a newly-installed MSCS cluster, you might find an error with Source COM+ and Event ID 4691 reported in the Application Event log.

This means that you are trying to run IBM MQ on a Microsoft Cluster Server (MSCS) environment when the Microsoft Distributed Transaction Coordinator (MSDTC) has not been configured to run in such an environment. For information on configuring MSDTC in a clustered environment, refer to Microsoft documentation.

### **IBM MQ MSCS support utility programs:**

A list of the IBM MQ support for MSCS utility programs that you can run at a command prompt.

IBM MQ support for MSCS includes the following utility programs:

#### **Register/unregister the resource type**

haregtyp.exe

After you *unregister* the IBM MQ MSCS resource type you can no longer create any resources of that type. MSCS does not let you unregister a resource type if you still have instances of that type within the cluster:

1. Using the MSCS Cluster Administrator, stop any queue managers that are running under MSCS control, by taking them offline as described in "Taking a queue manager offline from MSCS" on page 1070.
2. Using the MSCS Cluster Administrator, delete the resource instances.
3. At a command prompt, unregister the resource type by entering the following command:  
haregtyp /u

If you want to *register* the type (or re-register it at a later time), enter the following command at a command prompt:

haregtyp /r

After successfully registering the MSCS libraries, you must reboot the system if you have not done so since installing IBM MQ.

#### **Move a queue manager to MSCS storage**

hamvmqm.exe

See "Moving a queue manager to MSCS storage" on page 1063.

#### **Delete a queue manager from a node**

hadl1tmqm.exe

Consider the case where you have had a queue manager in your cluster, it has been moved from one node to another, and now you want to destroy it. Use the IBM MQ Explorer to delete it on

the node where it currently is. The registry entries for it still exist on the other computer. To delete these, enter the following command at a prompt on that computer:

```
hadltmqm /m qmname
```

where `qmname` is the name of the queue manager to remove.

### Check and save setup details

```
amqmsysn.exe
```

This utility presents a dialog showing full details of your IBM MQ MSCS Support setup, such as might be requested if you call IBM support. There is an option to save the details to a file.

## Multi-instance queue managers

Multi-instance queue managers are instances of the same queue manager configured on different servers. One instance of the queue manager is defined as the active instance and another instance is defined as the standby instance. If the active instance fails, the multi-instance queue manager restarts automatically on the standby server.

### Example multi-instance queue manager configuration

Figure 154 shows an example of a multi-instance configuration for queue manager QM1. IBM MQ is installed on two servers, one of which is a spare. One queue manager, QM1, has been created. One instance of QM1 is active, and is running on one server. The other instance of QM1 is running in standby on the other server, doing no active processing, but ready to take over from the active instance of QM1, if the active instance fails.




Figure 154. Multi-instance queue manager

When you intend to use a queue manager as a multi-instance queue manager, create a single queue manager on one of the servers using the `crtmqm` command, placing its queue manager data and logs in shared network storage. On the other server, rather than create the queue manager again, use the `addmqinf` command to create a reference to the queue manager data and logs on the network storage.

You can now run the queue manager from either of the servers. Each of the servers references the same queue manager data and logs; there is only one queue manager, and it is active on only one server at a time.

The queue manager can run either as a single instance queue manager, or as a multi-instance queue manager. In both cases only one instance of the queue manager is running, processing requests. The difference is that when running as a multi-instance queue manager, the server that is not running the active instance of the queue manager runs as a standby instance, ready to take over from the active instance automatically if the active server fails.

The only control you have over which instance becomes active first is the order in which you start the queue manager on the two servers. The first instance to acquire read/write locks to the queue manager data becomes the active instance.

You can swap the active instance to the other server, once it has started, by stopping the active instance using the switchover option to transfer control to the standby.

The active instance of QM1 has exclusive access to the shared queue manager data and logs folders when it is running. The standby instance of QM1 detects when the active instance has failed, and becomes the active instance. It takes over the QM1 data and logs in the state they were left by the active instance, and accepts reconnections from clients and channels.

The active instance might fail for various reasons that result in the standby taking over:

- Failure of the server hosting the active queue manager instance.
- Failure of connectivity between the server hosting the active queue manager instance and the file system.
- Unresponsiveness of queue manager processes, detected by IBM MQ, which then shuts down the queue manager.

You can add the queue manager configuration information to multiple servers, and choose any two servers to run as the active/standby pair. There is a limit of a total of two instances. You cannot have two standby instances and one active instance.

### **Additional components needed to build a high availability solution**

A multi-instance queue manager is one part of a high availability solution. You need some additional components to build a useful high availability solution.

- Client and channel reconnection to transfer IBM MQ connections to the computer that takes over running the active queue manager instance.
- A high performance shared network file system (NFS) that manages locks correctly and provides protection against media and file server failure.

**Important:** You must stop all multi-instance queue manager instances that are running in your environment before you can perform maintenance on the NFS drive. Make sure that you have queue manager configuration backups to recover, in the event of an NFS failure.

- Resilient networks and power supplies to eliminate single points of failure in the basic infrastructure.
- Applications that tolerate failover. In particular you need to pay close attention to the behavior of transactional applications, and to applications that browse IBM MQ queues.
- Monitoring and management of the active and standby instances to ensure that they are running, and to restart active instances that have failed. Although multi-instance queue managers restart automatically, you need to be sure that your standby instances are running, ready to take over, and that failed instances are brought back online as new standby instances.

IBM MQ MQI clients and channels reconnect automatically to the standby queue manager when it becomes active. More information about reconnection, and the other components in a high availability solution can be found in related topics. Automatic client reconnect is not supported by IBM MQ classes for Java.

## Supported platforms

You can create a multi-instance queue manager on any of the non- z/OS platforms supported by IBM MQ from IBM WebSphere MQ Version 7.0.1.

Automatic client reconnection is supported for MQI clients from IBM WebSphere MQ Version 7.0.1 onwards.

### Create a multi-instance queue manager:

Create a multi-instance queue manager, creating the queue manager on one server, and configuring IBM MQ on another server. Multi-instance queue managers share queue manager data and logs.

Most of the effort involved in creating a multi-instance queue manager is the task of setting up the shared queue manager data and log files. You must create shared directories on network storage, and make the directories available to other servers using network shares. These tasks need to be performed by someone with administrative authority, such as *root* on UNIX and Linux systems. The steps are as follows:

1. Create the shares for the data and log files.
2. Create the queue manager on one server.
3. Run the command **dspmqiinf** on the first server to collect the queue manager configuration data and copy it into the clipboard.
4. Run the command **addmqinf** with the copied data to create the queue manager configuration on the second server.

You do not run **crtmqm** to create the queue manager again on the second server.

### File access control

You need to take care that the user and group **mqm** on all other servers have permission to access the shares.

On UNIX and Linux, you need to make the **uid** and **gid** of **mqm** the same on all the systems. You might need to edit `/etc/passwd` on each system to set a common **uid** and **gid** for **mqm**, and then reboot your system.

On Microsoft Windows, the user ID that is running the queue manager processes must have full control permission to the directories containing the queue manager data and log files. You can configure the permission in two ways:

1. Create a queue manager with a global group as the alternative security principal. Authorize the global group to have full control access to the directories containing queue manager data and log files; see "Secure shared queue manager data and log directories and files on Windows" on page 1103. Make the user ID that is running the queue manager a member of the global group. You cannot make a local user a member of a global group, so the queue manager processes must run under a domain user ID. The domain user ID must be a member of the local group **mqm**. The task, "Create a multi-instance queue manager on domain workstations or servers" on page 1080, demonstrates how to set up a multi-instance queue manager using files secured in this way.
2. Create a queue manager on the domain controller, so that the local **mqm** group has domain scope, "domain local". Secure the file share with the domain local **mqm**, and run queue manager processes on all instances of a queue manager under the same domain local **mqm** group. The task, "Create a multi-instance queue manager on domain controllers" on page 1094, demonstrates how to set up a multi-instance queue manager using files secured in this way.

## Configuration information


Configure as many queue manager instances as you need by modifying the IBM MQ queue manager configuration information about each server. Each server must have the same version of IBM MQ installed at a compatible fix level. The commands, **dspmqinf** and **addmqinf** assist you to configure the additional queue manager instances. Alternatively, you can edit the `mqs.ini` and `qm.ini` files directly. The topics, “Create a multi-instance queue manager on Linux” on page 1115, “Create a multi-instance queue manager on domain workstations or servers” on page 1080, and “Create a multi-instance queue manager on domain controllers” on page 1094 are examples showing how to configure a multi-instance queue manager.


On Windows, UNIX and Linux systems, you can share a single `mqs.ini` file by placing it on the network share and setting the **AMQ\_MQS\_INI\_LOCATION** environment variable to point to it.


## Restrictions

1. Configure multiple instances of the same queue manager only on servers having the same operating system, architecture and endianness. For example, both machines must be either 32-bit or 64-bit.
2. All IBM MQ installations must be at release level 7.0.1 or higher.
3. Typically, active and standby installations are maintained at the same maintenance level. Consult the maintenance instructions for each upgrade to check whether you must upgrade all installations together.

Note that the maintenance levels for the active and passive queue managers must be identical.

4. Share queue manager data and logs only between queue managers that are configured with the same IBM MQ user, group, and access control mechanism.  For example, the network share set up on a Linux server could contain separate queue manager data and logs for UNIX and Linux queue managers, but could not contain the queue manager data used by IBM i.

 You can create multiple shares on the same networked storage for IBM i and other UNIX systems as long as the shares are different. You can give different shares different owners. The restriction is a consequence of the different names used for the IBM MQ users and groups between UNIX and IBM i. The fact that the user and group can have the same `uid` and `gid` does not relax the restriction.

5. On UNIX and Linux systems, configure the shared file system on networked storage with a hard, interruptible, mount rather than a soft mount. A hard interruptible mount forces the queue manager to hang until it is interrupted by a system call. Soft mounts do not guarantee data consistency after a server failure.
6. The shared log and data directories cannot be stored on a FAT, or an NFSv3 file system. For multi-instance queue managers on Windows, the networked storage must be accessed by the Common Internet File System (CIFS) protocol used by Windows networks.
7.  z/OS does not support multi-instance queue managers. Use queue sharing groups. Reconnectable clients do work with z/OS queue managers.

### *Windows domains and multi-instance queue managers:*

A multi-instance queue manager on Windows requires its data and logs to be shared. The share must be accessible to all instances of the queue manager running on different servers or workstations. Configure the queue managers and share as part of a Windows domain. The queue manager can run on a domain workstation or server, or on the domain controller.

Before configuring a multi-instance queue manager, read “Secure unshared queue manager data and log directories and files on Windows” on page 1106 and “Secure shared queue manager data and log directories and files on Windows” on page 1103 to review how to control access to queue manager data and log files. The topics are educational; if you want to go directly to setting up shared directories for a multi-instance queue manager in a Windows domain; see “Create a multi-instance queue manager on domain workstations or servers” on page 1080.

### **Run a multi-instance queue manager on domain workstations or servers**

From Version 7.1, multi-instance queue managers run on a workstation or server that is a member of a domain. Before Version 7.1, multi-instance queue managers ran only on domain controllers; see “Run a multi-instance queue manager on domain controllers” on page 1080. To run a multi-instance queue manager on Windows, you require a domain controller, a file server, and two workstations or servers running the same queue manager connected to the same domain.

The change that makes it possible to run a multi-instance queue manager on any server or workstation in a domain, is that you can now create a queue manager with an additional security group. The additional security group is passed in the `crtmqm` command, in the `-a` parameter. You secure the directories that contain the queue manager data and logs with the group. The user ID that runs queue manager processes must be a member of this group. When the queue manager accesses the directories, Windows checks the permissions the user ID has to access the directories. By giving both the group and the user ID domain scope, the user ID running the queue manager processes has credentials from the global group. When the queue manager is running on a different server, the user ID running the queue manager processes can have the same credentials. The user ID does not have to be the same. It has to be a member of the alternative security group, as well as a member of the local `mqm` group.

The task of creating a multi-instance queue manager is the same as in Version 7.0.1 with one change. You must add the additional security group name to the parameters of the `crtmqm` command. The task is described in “Create a multi-instance queue manager on domain workstations or servers” on page 1080.

Multiple steps are required to configure the domain, and the domain servers and workstations. You must understand how Windows authorizes access by a queue manager to its data and log directories. If you are not sure how queue manager processes are authorized to access their log and data files read the topic “Secure unshared queue manager data and log directories and files on Windows” on page 1106. The topic includes two tasks to help you understand the steps the required. The tasks are “Reading and writing data and log files authorized by the local `mqm` group” on page 1108 and “Reading and writing data and log files authorized by an alternative local security group” on page 1111. Another topic, “Secure shared queue manager data and log directories and files on Windows” on page 1103, explains how to secure shared directories containing queue manager data and log files with the alternative security group. The topic includes four tasks, to set up a Windows domain, create a file share, install IBM MQ for Windows, and configure a queue manager to use the share. The tasks are as follows:

1. “Creating an Active Directory and DNS domain for IBM MQ” on page 1083.
2. “Installing IBM MQ on a server or workstation in a Windows domain” on page 1086.
3. “Creating a shared directory for queue manager data and log files” on page 1089.
4. “Reading and writing shared data and log files authorized by an alternative global security group” on page 1091.

You can then do the task, “Create a multi-instance queue manager on domain workstations or servers,” using the domain. Do these tasks to explore setting up a multi-instance queue manager before transferring your knowledge to a production domain.

### Run a multi-instance queue manager on domain controllers

In Version 7.0.1, multi-instance queue managers ran only on domain controllers. Queue manager data could be secured with the domain `mqm` group. As the topic “Secure shared queue manager data and log directories and files on Windows” on page 1103 explains, you cannot share directories secured with the local `mqm` group on workstations or servers. However on domain controllers all group and principals have domain scope. If you install IBM MQ for Windows on a domain controller, the queue manager data and log files are secured with the domain `mqm` group, which can be shared. Follow the steps in the task, “Create a multi-instance queue manager on domain controllers” on page 1094 to configure a multi-instance queue manager on domain controllers.

#### Related information:

[Managing Authorization and Access Control](#)

[How to use Windows Server cluster nodes as domain controllers](#)

*Create a multi-instance queue manager on domain workstations or servers:*

An example shows how to set up a multi-instance queue manager on Windows on a workstation or a server that is part of a Windows domain. The server does not have to be a domain controller. The setup demonstrates the concepts involved, rather than being production scale. The example is based on Windows Server 2008. The steps might differ on other versions of Windows Server.

In a production scale configuration, you might have to tailor the configuration to an existing domain. For example, you might define different domain groups to authorize different shares, and to group the user IDs that run queue managers.

The example configuration consists of three servers:

- sun* A Windows Server 2008 domain controller. It owns the *wmq.example.com* domain that contains *Sun* , *mars* , and *venus* . For the purposes of illustration, it is also used as the file server.
- mars* A Windows Server 2008 used as the first IBM MQ server. It contains one instance of the multi-instance queue manager called *QMGR* .
- venus* A Windows Server 2008 used as the second IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .

Replace the italicized names in the example, with names of your choosing.

#### Before you begin

On Windows, you do not need to verify the file system that you plan to store queue manager data and log files on. The checking procedure, Verifying shared file system behavior, is applicable to UNIX and Linux. On Windows, the checks are always successful.

Do the steps in the following tasks. The tasks create the domain controller and domain, install IBM MQ for Windows on one server, and create the file share for data and log files. If you are configuring an existing domain controller, you might find it useful to try out the steps on a new Windows Server 2008. You can adapt the steps to your domain.

1. “Creating an Active Directory and DNS domain for IBM MQ” on page 1083.
2. “Installing IBM MQ on a server or workstation in a Windows domain” on page 1086.
3. “Creating a shared directory for queue manager data and log files” on page 1089.



4. "Reading and writing shared data and log files authorized by an alternative global security group" on page 1091.

### About this task

This task is one of a sequence of tasks to configure a domain controller and two servers in the domain to run instances of a queue manager. In this task you configure a second server, *venus*, to run another instance of the queue manager *QMGR*. Follow the steps in this task to create the second instance of the queue manager, *QMGR*, and test that it works.

This task is separate from the four tasks in the preceding section. It contains the steps that convert a single instance queue manager into a multi-instance queue manager. All the other steps are common to single or multi-instance queue managers.

### Procedure

1. Configure a second server to run IBM MQ for Windows.
  - a. Do the steps in the task "Installing IBM MQ on a server or workstation in a Windows domain" on page 1086 to create a second domain server. In this sequence of tasks the second server is called *venus*.

**Tip:** Create the second installation using the same installation defaults for IBM MQ on each of the two servers. If the defaults differ, you might have to tailor the `Prefix` and the `InstallationName` variables in the *QMGR QueueManager* stanza in the IBM MQ configuration file `mqs.ini`. The variables refer to paths that can differ for each installation and queue manager on each server. If the paths remain the same on every server, it is simpler to configure a multi-instance queue manager.

2. Create a second instance of *QMGR* on *venus*.
  - a. If *QMGR* on *mars* does not exist, do the task "Reading and writing shared data and log files authorized by an alternative global security group" on page 1091, to create it
  - b. Check the values of the `Prefix` and `InstallationName` parameters are correct for *venus*.

On *mars*, run the **dspmqrinf** command:

```
dspmqrinf QMGR
```

The system response:

```
QueueManager:
Name=QMGR
Directory=QMGR
Prefix=C:\ProgramData\IBM\MQ
DataPath=\\sun\wmq\data\QMGR
InstallationName=Installation1
```

- c. Copy the machine-readable form of the *QueueManager* stanza to the clipboard. On *mars* run the **dspmqrinf** command again, with the **-o command** parameter.

```
dspmqrinf -o command QMGR
```

The system response:

```
addmqinf -s QueueManager -v Name=QMGR
-v Directory=QMGR -v Prefix="C:\ProgramData\IBM\MQ"
-v DataPath=\\sun\wmq\data\QMGR
```

- d. On *venus* run the **addmqinf** command from the clipboard to create an instance of the queue manager on *venus*.

Adjust the command, if necessary, to accommodate differences in the `Prefix` or `InstallationName` parameters.

```
addmqinf -s QueueManager -v Name=QMGR
-v Directory=QMGR -v Prefix="C:\ProgramData\IBM\MQ"
-v DataPath=\\sun\wmq\data\QMGR
```

IBM MQ configuration information added.

3. Start the queue manager *QMGR* on *venus* , permitting standby instances.
  - a. Check *QMGR* on *mars* is stopped.

On *mars* , run the **dspmq** command:

```
dspmq -m QMGR
```

The system response depends on how the queue manager was stopped; for example:

```
C:\Users\Administrator>dspmq -m QMGR
QMNAME(QMGR) STATUS(Ended immediately)
```

- b. On *venus* run the **strmqm** command to start *QMGR* permitting standbys:

```
strmqm -x QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation 'Installation1'.
5 log records accessed on queue manager 'QMGR' during the log
replay phase.
Log replay for queue manager 'QMGR' complete.
Transaction manager state recovered for queue manager 'QMGR'.
IBM MQ queue manager 'QMGR' started using V7.1.0.0.
```

## Results

To test the multi-instance queue manager switches over, do the following steps:

1. On *mars* , run the **strmqm** command to start *QMGR* permitting standbys:

```
strmqm -x QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation 'Installation1'.
A standby instance of queue manager 'QMGR' has been started.
The active instance is running elsewhere.
```

2. On *venus* run the **endmqm** command:

```
endmqm -r -s -i QMGR
```

The system response on *venus* :

```
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ended, permitting switchover to
a standby instance.
```

And on *mars* :

```
dspmqr
QMNAME(QMGR) STATUS(Running as standby)
C:\Users\wmquser2>dspmqr
QMNAME(QMGR) STATUS(Running as standby)
C:\Users\wmquser2>dspmqr
QMNAME(QMGR) STATUS(Running)
```

## What to do next

To verify a multi-instance queue manager using sample programs; see “Verify the multi-instance queue manager on Windows” on page 1101.

*Creating an Active Directory and DNS domain for IBM MQ:*

This task creates the domain *wmq.example.com* on a Windows 2008 domain controller called *sun*. It configures the Domain *mqm* global group in the domain, with the correct rights, and with one user.

In a production scale configuration, you might have to tailor the configuration to an existing domain. For example, you might define different domain groups to authorize different shares, and to group the user IDs that run queue managers.

The example configuration consists of three servers:

- sun* A Windows Server 2008 domain controller. It owns the *wmq.example.com* domain that contains *Sun*, *mars*, and *venus*. For the purposes of illustration, it is also used as the file server.
- mars* A Windows Server 2008 used as the first IBM MQ server. It contains one instance of the multi-instance queue manager called *QMGR*.
- venus* A Windows Server 2008 used as the second IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR*.

Replace the italicized names in the example, with names of your choosing.

## Before you begin

1. The task steps are consistent with a Windows Server 2008 that is installed but not configured with any roles. If you are configuring an existing domain controller, you might find it useful to try out the steps on a new Windows Server 2008. You can adapt the steps to your domain.

## About this task

In this task, you create an Active Directory and DNS domain on a new domain controller. You then configure it ready to install IBM MQ on other servers and workstations that join the domain. Follow the task if you are unfamiliar with installing and configuring Active Directory to create a Windows domain. You must create a Windows domain in order to create a multi-instance queue manager configuration. The task is not intended to guide you in the best way to configure a Windows domain. To deploy multi-instance queue managers in a production environment, you must consult Windows documentation.

During the task you do the following steps:

1. Install Active Directory.
2. Add a domain.
3. Add the domain to DNS.
4. Create the global group Domain *mqm* and give it the correct rights.
5. Add a user and make it a member of the global group Domain *mqm*.

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

For the purposes of the task the domain controller hostname is *sun* , and the two IBM MQ servers are called *mars* and *venus* . The domain is called *wmq.example.com* . You can replace all the italicized names in the task with names of your own choosing.

### Procedure

1. Log on to the domain controller, *sun* , as the local or Workgroup administrator.

If the server is already configured as a domain controller, you must log on as a domain administrator.

2. Run the Active Directory Domain Services wizard.

- a. Click **Start > Run...** Type `dcpromo` and click **OK**.

If the Active Directory binary files are not already installed, Windows installs the files automatically.

3. In the first window of the wizard, leave the **Use advanced mode installation** check box clear. Click **Next > Next** and click **Create a new domain in a new forest > Next**.

4. Type *wmq.example.com* into the **FQDN of the forest root domain** field. Click **Next**.

5. In the Set Forest Functional Level window, select **Windows Server 2003**, or later, from the list of **Forest functional levels > Next**.

The oldest level of Windows Server that is supported by IBM MQ is Windows Server 2003.

6. Optional: In the Set Domain Functional Level window, select **Windows Server 2003**, or later, from the list of **Domain functional levels > Next**.

This step is only required if you set the Forest Functional Level to **Windows Server 2003**.

7. The Additional Domain Controller Options window opens, with **DNS server** selected as an additional option. Click **Next** and **Yes** to clear the warning window.

**Tip:** If a DNS server is already installed this option is not presented to you. If you want to follow this task precisely, remove all the roles from this domain controller and start again.

8. Leave the Database, Log Files, and SYSVOL directories unchanged; click **Next**.

9. Type a password into the **Password** and **Confirm password** fields in the Directory Services Restore Mode Administrator Password window. Click **Next > Next**. Select **Reboot on completion** in the final wizard window.

10. When the domain controller reboots, log on as *wmq\Administrator*.

The server manager starts automatically.

11. Open the *wmq.example.com\Users* folder

- a. Open **Server Manager > Roles > Active Directory Domain Services > *wmq.example.com* > Users**.

12. Right-click **Users > New > Group**.

- a. Type a group name into the **Group name** field.

**Note:** The preferred group name is Domain `mqm`. Type it exactly as shown.

- Calling the group Domain `mqm` modifies the behavior of the “Prepare IBM MQ ” wizard on a domain workstation or server. It causes the “Prepare IBM MQ ” wizard automatically to add the group Domain `mqm` to the local `mqm` group on each new installation of IBM MQ in the domain.
- You can install workstations or servers in a domain with no Domain `mqm` global group. If you do so, you must define a group with the same properties as Domain `mqm` group. You must make that group, or the users that are members of it, members of the local `mqm` group wherever IBM

MQ is installed in a domain. You can place domain users into multiple groups. Create multiple domain groups, each group corresponding to a set of installations that you want to manage separately. Split domain users, according to the installations they manage, into different domain groups. Add each domain group or groups to the local `mqm` group of different IBM MQ installations. Only domain users in the domain groups that are members of a specific local `mqm` group can create, administer, and run queue managers for that installation.

- The domain user that you nominate when installing IBM MQ on a workstation or server in a domain must be a member of the `Domain mqm` group, or of an alternative group you defined with same properties as the `Domain mqm` group.
  - b. Leave **Global** clicked as the **Group scope**, or change it to **Universal**. Leave **Security** clicked as the **Group type**. Click **OK**.
13. Add the rights, **Allow Read group membership** and **Allow Read groupMembershipSAM** to the rights of the `Domain mqm` global group.
- a. In the Server Manager action bar, click **View > Advanced features**
  - b. In the Server Manager navigation tree, click **Users**
  - c. In the Users window, right-click **Domain mqm > Properties**
  - d. Click **Security > Advanced > Add....** Type `Domain mqm` and click **Check names > OK**.  
The **Name** field is prefilled with the string, `Domain mqm (domain name\Domain mqm)`.
  - e. Click **Properties**. In the **Apply to** list, select **Descendant User Objects**.
  - f. From the **Permissions** list, select the **Read group membership** and **Read groupMembershipSAM** **Allow** check boxes; click **OK > Apply > OK > OK**.
14. Add two or more users to the `Domain mqm` global group.
- One user, `wmquser1` in the example, runs the IBM MQ service, and the other user, `wmquser2`, is used interactively.
- A domain user is required to create a queue manager that uses the alternative security group in a domain configuration. It is not sufficient for the user ID to be an administrator, although an administrator has authority to run the `crtmqm` command. The domain user, who could be an administrator, must be a member of the local `mqm` group as well as of the alternative security group.
- In the example, you make `wmquser1` and `wmquser2` members of the `Domain mqm` global group. The "Prepare IBM MQ" wizard automatically configures `Domain mqm` as a member of the local `mqm` group where ever the wizard is run.
- You must provide a different user to run the IBM MQ service for each installation of IBM MQ on a single computer. You can reuse the same users on different computers.
- a. In the Server Manager navigation tree, click **Users > New > User**
  - b. In the New Object - User window, type `wmquser1` into the **User logon name** field. Type `WebSphere` into the **First name** field, and `MQ1` into the **Last name** field. Click **Next**.
  - c. Type a password into the **Password** and **Confirm password** fields, and clear the **User must change password at next logon** check box. Click **Next > Finish**.
  - d. In the Users window, right-click `WebSphere MQ > Add to a group....` Type `Domain mqm` and click **Check Names > OK > OK**.
  - e. Repeat steps a to d to add `WebSphere MQ2` as `wmquser2`.
15. Running IBM MQ as a service. If you need to run IBM MQ as a service, and then give the domain user (that you obtained from your domain administrator) the access to run as a service, carry out the following procedure:
- a. Click **Start > Run....** Type the command `secpol.msc` and click **OK**.
  - b. Open **Security Settings > Local Policies > User Rights Assignments**. In the list of policies, right-click **Log on as a service > Properties**.
  - c. Click **Add User or Group...** Type the name of the user you obtained from your domain administrator, and click **Check Names**

- d. If prompted by a Windows Security window, type the user name and password of an account user or administrator with sufficient authority, and click **OK > Apply > OK**. Close the Local Security Policy window.

**Note:** On Windows Server 2008 and Windows Server 2012 the User Account Control (UAC) is enabled by default.

The UAC feature restricts the actions users can perform on certain operating system facilities, even if they are members of the Administrators group. You must take appropriate steps to overcome this restriction.

### What to do next

Proceed to the next task, “Installing IBM MQ on a server or workstation in a Windows domain.”

*Installing IBM MQ on a server or workstation in a Windows domain:*

In this task, you install and configure IBM MQ on a server or workstation in the *wmq.example.com* Windows domain.

In a production scale configuration, you might have to tailor the configuration to an existing domain. For example, you might define different domain groups to authorize different shares, and to group the user IDs that run queue managers.

The example configuration consists of three servers:

- sun* A Windows Server 2008 domain controller. It owns the *wmq.example.com* domain that contains *Sun*, *mars*, and *venus*. For the purposes of illustration, it is also used as the file server.
- mars* A Windows Server 2008 used as the first IBM MQ server. It contains one instance of the multi-instance queue manager called *QMGR*.
- venus* A Windows Server 2008 used as the second IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR*.

Replace the italicized names in the example, with names of your choosing.

### Before you begin

1. Do the steps in “Creating an Active Directory and DNS domain for IBM MQ” on page 1083 to create a domain controller, *sun*, for the domain *wmq.example.com*. Change the italicized names to suit your configuration.
2. See Hardware and software requirements on Windows systems for other Windows versions you can run IBM MQ on.

### About this task

In this task you configure a Windows Server 2008, called *mars*, as a member of the *wmq.example.com* domain. You install IBM MQ, and configure the installation to run as a member of the *wmq.example.com* domain.

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

For the purposes of the task the domain controller hostname is *sun* , and the two IBM MQ servers are called *mars* and *venus* . The domain is called *wmq.example.com* . You can replace all the italicized names in the task with names of your own choosing.

### Procedure

1. Add the domain controller, *sun.wmq.example.com* to *mars* as a DNS server.
  - a. On *mars* , log on as *mars\Administrator* and click **Start**.
  - b. Right-click **Network > Properties > Manage network connections**.
  - c. Right-click the network adapter, click **Properties**.

The system responds with the Local Area Connection Properties window listing items the connection uses.
  - d. Select the **Internet Protocol Version 4** or **Internet Protocol Version 6** from the list of items in the Local Area Connection Properties window. Click **Properties > Advanced...** and click the DNS tab.
  - e. Under the DNS server addresses, click **Add...**
  - f. Type the IP address of the domain controller, which is also the DNS server, and click **Add**.
  - g. Click **Append these DNS suffixes > Add...**
  - h. Type *wmq.example.com* and click **Add**.
  - i. Type *wmq.example.com* in the **DNS suffix for this connection** field.
  - j. Select **Register this connection's address in DNS** and **Use this connection's suffix in DNS registration**. Click **OK > OK > Close**
  - k. Open a command window, and type the command **ipconfig /all** to review the TCP/IP settings.
2. On *mars* , add the computer to the *wmq.example.com* domain.
  - a. Click **Start**
  - b. Right-click **Computer > Properties**. In the Computer name, domain and workgroup settings division, click **Change settings**.
  - c. In the System Properties windows, click **Change...**
  - d. Click **Domain**, type *wmq.example.com* , and click **OK**.
  - e. Type the **User name** and **Password** of the domain controller administrator, who has the authority to permit the computer to join the domain, and click **OK**.
  - f. Click **OK > OK > Close > Restart Now** in response to the "Welcome to the *wmq.example.com* domain" message.
3. Check that the computer is a member of the *wmq.example.com* domain
  - a. On *sun* , log on to the domain controller as *wmq\Administrator*.
  - b. Open **Server Manager > Active Directory Domain Services > wmq.example.com > Computers** and check *mars* is listed correctly in the Computers window.
4. Install IBM MQ for Windows on *mars* .

For further information about running the IBM MQ for Windows installation wizard; see Installing IBM MQ server on Windows .

  - a. On *mars* , log on as the local administrator, *mars\Administrator*.
  - b. Run the **Setup** command on the IBM MQ for Windows installation media.

The IBM MQ Launchpad application starts.
  - c. Click **Software Requirements** to check that the prerequisite software is installed.
  - d. Click **Network Configuration > Yes** to configure a domain user ID.

The task, "Creating an Active Directory and DNS domain for IBM MQ" on page 1083, configures a domain user ID for this set of tasks.
  - e. Click **IBM MQ Installation**, select an installation language and click **Launch IBM MQ Installer**.
  - f. Confirm the license agreement and click **Next > Next > Install** to accept the default configuration. Wait for the installation to complete, and click **Finish**.

You might want to change the name of the installation, install different components, configure a different directory for queue manager data and logs, or install into a different directory. If so, click **Custom** rather than **Typical**. IBM MQ is installed, and the installer starts the “Prepare IBM MQ ” wizard.

**Important:** Do not run the wizard yet.

5. Configure the user that is going to run the IBM MQ service with the **Run as a service** right. Choose whether to configure the local `mqm` group, the Domain `mqm` group, or the user that is going to run the IBM MQ service with the right. In the example, you give the user the right.
  - a. Click **Start > Run...**, type the command `secpol.msc` and click **OK**.
  - b. Open **Security Settings > Local Policies > User Rights Assignments**. In the list of policies, right-click **Log on as a service > Properties** .
  - c. Click **Add User or Group...** and type `wmquser1` and click **Check Names**
  - d. Type the user name and password of a domain administrator, `wmq\Administrator`, and click **OK > Apply > OK**. Close the Local Security Policy window.
6. Run the “Prepare IBM MQ ” wizard.

For further information about running the “Prepare IBM MQ ” wizard; see Configuring IBM MQ with the Prepare IBM MQ wizard.

  - a. The IBM MQ Installer runs the “Prepare IBM MQ ” automatically.

To start the wizard manually, find the shortcut to the “Prepare IBM MQ ” in the **Start > All programs > IBM MQ** folder. Select the shortcut that corresponds to the installation of IBM MQ in a multi-installation configuration.
  - b. Click **Next** and leave **Yes** clicked in response to the question “Identify if there is a Windows 2000 or later domain controller in the network”.
  - c. Click **Yes > Next** in the first Configuring IBM MQ for Windows for Windows domain users window.
  - d. In the second Configuring IBM MQ for Windows for Windows domain users window, type `wmq` in the **Domain** field. Type `wmquser1` in the **User name** field, and the password, if you set one, in the **Password** field. Click **Next**.

The wizard configures and starts the IBM MQ with `wmquser1`.
  - e. In the final page of the wizard, select or clear the check boxes as you require and click **Finish**.

#### What to do next

1. Do the task, “Reading and writing data and log files authorized by the local `mqm` group” on page 1108, to verify that the installation and configuration are working correctly.
2. Do the task, “Creating a shared directory for queue manager data and log files” on page 1089, to configure a file share to store the data and log files of a multi-instance queue manager.



## Related information:

User rights required for an IBM MQ Windows Service

*Creating a shared directory for queue manager data and log files:*

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing.

In a production scale configuration, you might have to tailor the configuration to an existing domain. For example, you might define different domain groups to authorize different shares, and to group the user IDs that run queue managers.

The example configuration consists of three servers:

- sun* A Windows Server 2008 domain controller. It owns the *wmq.example.com* domain that contains *Sun* , *mars* , and *venus* . For the purposes of illustration, it is also used as the file server.
- mars* A Windows Server 2008 used as the first IBM MQ server. It contains one instance of the multi-instance queue manager called *QMGR* .
- venus* A Windows Server 2008 used as the second IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .

Replace the italicized names in the example, with names of your choosing.

## Before you begin

1. To do this task exactly as documented, do the steps in the task, “Creating an Active Directory and DNS domain for IBM MQ” on page 1083, to create the domain *sun.wmq.example.com* on the domain controller *sun* . Change the italicized names to suit your configuration.

## About this task

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

In the task, you create a share containing a data and log directory, and a global group to authorize access to the share. You pass the name of the global group that authorizes the share to the **crtmqm** command in its **-a** parameter. The global group gives you the flexibility of separating the users of this share from users of other shares. If you do not need this flexibility, authorize the share with the Domain **mqm** group rather than create a new global group.

The global group used for sharing in this task is called *wmqha* , and the share is called *wmq* . They are defined on the domain controller *sun* in the Windows domain *wmq.example.com* . The share has full control permissions for the global group *wmqha* . Replace the italicized names in the task with names of your choosing.

For the purposes of this task the domain controller is the same server as the file server. In practical applications, split the directory and file services between different servers for performance and availability.

You must configure the user ID that the queue manager is running under to be a member of two groups. It must be a member of the local **mqm** group on an IBM MQ server, and of the *wmqha* global group.

In this set of tasks, when the queue manager is running as a service, it runs under the user ID *wmquser1* , so *wmquser1* must be a member of *wmqha* . When the queue manager is running interactively, it runs under the user ID *wmquser2* , so *wmquser2* must be a member of *wmqha* . Both *wmquser1* and *wmquser2* are members of the global group Domain *mqm*. Domain *mqm* is a member of the local *mqm* group on the *mars* and *venus* IBM MQ servers. Hence, *wmquser1* and *wmquser2* are members of the local *mqm* group on both IBM MQ servers.

### Procedure

1. Log on to the domain controller, *sun.wmq.example.com* as the domain administrator.
2. Create the global group *wmqha* .
  - a. Open **Server Manager** > **Roles** > **Active Directory Domain Services** > *wmq.example.com* > **Users**.
  - b. Open the *wmq.example.com*\Users folder
  - c. Right-click **Users** > **New** > **Group**.
  - d. Type *wmqha* into the **Group name** field.
  - e. Leave **Global** clicked as the **Group scope** and **Security** as the **Group type**. Click **OK**.
3. Add the domain users *wmquser1* and *wmquser2* to the global group, *wmqha* .
  - a. In the Server Manager navigation tree, click **Users** and right-click *wmqha* > **Properties** in the list of users.
  - b. Click the Members tab in the *wmqha* Properties window.
  - c. Click **Add...** ; type *wmquser1* ; *wmquser2* and click **Check Names** > **OK** > **Apply** > **OK**.
4. Create the directory tree to contain queue manager data and log files.
  - a. Open a command prompt.
  - b. Type the command:

```
md c:\wmq\data, c:\wmq\logs
```
5. Authorize the global group *wmqha* to have full control permission to the *c:\wmq* directories and share.
  - a. In Windows Explorer, right-click *c:\wmq* > **Properties**.
  - b. Click the **Security** tab and click **Advanced** > **Edit...**
  - c. Clear the check box for **Include inheritable permissions from this object's owner**. Click **Copy** in the Windows Security window.
  - d. Select the lines for Users in the list of **Permission entries** and click **Remove**. Leave the lines for SYSTEM, Administrators, and CREATOR OWNER in the list of **Permission entries**.
  - e. Click **Add...**, and type the name of the global group *wmqha* . Click **Check Names** > **OK**.
  - f. In the Permission Entry for *wmq* window, select **Full Control** in the list of **Permissions**.
  - g. Click **OK** > **Apply** > **OK** > **OK** > **OK**
  - h. In Windows Explorer, right-click *c:\wmq* > **Share...**
  - i. Click **Advanced Sharing...** and select the **Share this folder** check box. Leave the share name as *wmq* .
  - j. Click **Permissions** > **Add...**, and type the name of the global group *wmqha* . Click **Check Names** > **OK**.
  - k. Select *wmqha* in the list of **Group or user names**. Select the **Full Control** check box in the list of **Permissions for wmqha** ; click **Apply**.
  - l. Select *Administrators* in the list of **Group or user names**. Select the **Full Control** check box in the list of **Permissions for Administrators** ; click **Apply** > **OK** > **OK** > **Close**.

### What to do next

Check that you can read and write files to the shared directories from each of the IBM MQ servers. Check the IBM MQ service user ID, *wmquser1* and the interactive user ID, *wmquser2*.

1. If you are using remote desktop, you must add *wmq\wmquser1* and *wmquser2* to the local group Remote Desktop Users on *mars* .
  - a. Log on to *mars* as *wmq\Administrator*
  - b. Run the **lusrmgr.msc** command to open the Local Users and Groups window.
  - c. Click **Groups**. Right-click **Remote Desktop Users > Properties > Add...** Type *wmquser1* ; *wmquser2* and click **Check Names**.
  - d. Type in the user name and password of the domain administrator, *wmq\Administrator*, and click **OK > Apply > OK**.
  - e. Close the Local Users and Groups window.
2. Log on to *mars* as *wmq\wmquser1*.
  - a. Open a Windows Explorer window, and type in `\\sun\wmq` .  
The system responds by opening the *wmq* share on *sun.wmq.example.com* , and lists the data and logs directories.
  - b. Check the permissions of *wmquser1* by creating a file in data subdirectory, adding some content, reading it, and then deleting it.
3. Log on to *mars* as *wmq\wmquser2*, and repeat the checks.
4. Do the next task, to create a queue manager to use the shared data and log directories; see “Reading and writing shared data and log files authorized by an alternative global security group.”

*Reading and writing shared data and log files authorized by an alternative global security group:*

This task shows how to use the **-a** flag on the **crtmqm** command. The **-a** flag gives the queue manager access to its log and data files on a remote file share using the alternative security group.

In a production scale configuration, you might have to tailor the configuration to an existing domain. For example, you might define different domain groups to authorize different shares, and to group the user IDs that run queue managers.

The example configuration consists of three servers:

- sun* A Windows Server 2008 domain controller. It owns the *wmq.example.com* domain that contains *Sun* , *mars* , and *venus* . For the purposes of illustration, it is also used as the file server.
- mars* A Windows Server 2008 used as the first IBM MQ server. It contains one instance of the multi-instance queue manager called *QMGR* .
- venus* A Windows Server 2008 used as the second IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .

Replace the italicized names in the example, with names of your choosing.

### **Before you begin**

Do the steps in the following tasks. The tasks create the domain controller and domain, install IBM MQ for Windows on one server, and create the file share for data and log files. If you are configuring an existing domain controller, you might find it useful to try out the steps on a new Windows Server 2008. You can adapt the steps to your domain.

1. “Creating an Active Directory and DNS domain for IBM MQ” on page 1083.
2. “Installing IBM MQ on a server or workstation in a Windows domain” on page 1086.
3. “Creating a shared directory for queue manager data and log files” on page 1089.

## About this task

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, "Windows domains and multi-instance queue managers" on page 1079.

In this task, you create a queue manager that stores its data and logs in a remote directory on a file server. For the purposes of this example, the file server is the same server as the domain controller. The directory containing the data and log folders is shared with full control permission given to the global group `wmqha`.

## Procedure

1. Log on to the domain server, `mars`, as the local administrator, `mars\Administrator`.
2. Open a command window.
3. Restart the IBM MQ service.

You must restart the service so that the user ID it runs under acquires the additional security credentials you configured for it.

Type the commands:

```
endmqsvc
strmqsvc
```

The system responses:

```
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
The MQ service for installation 'Installation1' ended successfully.
```

And:

```
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
The MQ service for installation 'Installation1' started successfully.
```

4. Create the queue manager.

```
crtmqm -a wmq\wqha -sax -u SYSTEM.DEAD.LETTER.QUEUE -md \\sun\wmq\data -ld \\sun\wmq\logs QMGR
```

You must specify the domain, `wmq`, of the alternative security group `wmqha` by specifying full domain name of the global group "`wmq\wqha`".

You must spell out the Universal Naming Convention (UNC) name of the share `\\sun\wmq`, and not use a mapped drive reference.

The system response:

```
IBM MQ queue manager created.
Directory '\\sun\wmq\data\QMGR' created.
The queue manager is associated with installation '1'
Creating or replacing default objects for queue manager 'QMGR'
Default objects statistics : 74 created. 0 replaced.
Completing setup.
Setup completed.
```

## What to do next

Test the queue manager by putting and getting a message to a queue.

1. Start the queue manager.

```
strmqm QMGR
```

The system response:

IBM MQ queue manager 'QMGR' starting.  
The queue manager is associated with installation '1'.  
5 log records accessed on queue manager 'QMGR' during the log  
replay phase.  
Log replay for queue manager 'QMGR' complete.  
Transaction manager state recovered for queue manager 'QMGR'.  
IBM MQ queue manager 'QMGR' started using V7.1.0.0.

2. Create a test queue.

```
echo define qlocal(QTEST) | runmqsc QMGR
```

The system response:

```
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QMGR.
```

```
1 : define qlocal(QTEST)
AMQ8006: IBM MQ queue created.
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

3. Put a test message using the sample program **amqsput**.

```
echo 'A test message' | amqsput QTEST QMGR
```

The system response:

```
Sample AMQSPUT0 start
target queue is QTEST
Sample AMQSPUT0 end
```

4. Get the test message using the sample program **amqsget**.

```
amqsget QTEST QMGR
```

The system response:

```
Sample AMQSGET0 start
message <A test message>
Wait 15 seconds ...
no more messages
Sample AMQSGET0 end
```

5. Stop the queue manager.

```
endmqm -i QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ended.
```

6. Delete the queue manager.

```
dltmqm QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' deleted.
```

7. Delete the directories you created.

**Tip:** Add the /Q option to the commands to prevent the command prompting to delete each file or directory.

```
del /F /S C:\wmq*.*
rmdir /S C:\wmq
```

Create a multi-instance queue manager on domain controllers:

An example shows how to set up a multi-instance queue manager on Windows on domain controllers. The setup demonstrates the concepts involved, rather than being production scale. The example is based on Windows Server 2008. The steps might differ on other versions of Windows Server.

The configuration uses the concept of a mini-domain, or “domainlet” ; see Windows 2000, Windows Server 2003, and Windows Server 2008 cluster nodes as domain controllers. To add multi-instance queue managers to an existing domain, see “Create a multi-instance queue manager on domain workstations or servers” on page 1080.

The example configuration consists of three servers:

- sun* A Windows Server 2008 server used as the first domain controller. It defines the *wmq.example.com* domain that contains *sun* , *earth* , and *mars* . It contains one instance of the multi-instance queue manager called *QMGR* .
- earth* A Windows Server 2008 used as the second domain controller IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .
- mars* A Windows Server 2008 used as the file server.

Replace the italicized names in the example, with names of your choosing.

### Before you begin

1. On Windows, you do not need to verify the file system that you plan to store queue manager data and log files on. The checking procedure, Verifying shared file system behavior, is applicable to UNIX and Linux. On Windows, the checks are always successful.
2. Do the steps in “Creating an Active Directory and DNS domain for IBM MQ” on page 1083 to create the first domain controller.
3. Do the steps in “Adding a second domain controller to the *wmq.example.com* domain” on page 1097 to add a second domain controller, install IBM MQ for Windows on both domain controllers, and verify the installations.
4. Do the steps in “Installing IBM MQ on domain controllers in the *wmq.example.com* domain” on page 1099 to install IBM MQ on the two domain controllers.

### About this task

On a file server in the same domain create a share for the queue manager log and data directories. Next, create the first instance of a multi-instance queue manager that uses the file share on one of the domain controllers. Create the other instance on the other domain controller and finally verify the configuration. You can create the file share on a domain controller.

In the sample, *sun* is the first domain controller, *earth* the second, and *mars* is the file server.

### Procedure

1. Create the directories that are to contain the queue manager data and log files.
  - a. On *mars* , type the command:

```
md c:\wmq\data , c:\wmq\logs
```
2. Share the directories that are to contain the queue manager data and log files.

You must permit full control access to the domain local group *mqm*, and the user ID you use to create the queue manager. In the example, user IDs that are members of Domain Administrators have the authority to create queue managers.

The file share must be on a server that is in the same domain as the domain controllers. In the example, the server *mars* is in the same domain as the domain controllers.

- a. In Windows Explorer, right-click *c:\wmq* > **Properties**.
- b. Click the **Security** tab and click **Advanced** > **Edit...**
- c. Clear the check box for **Include inheritable permissions from this object's owner**. Click **Copy** in the Windows Security window.
- d. Select the lines for Users in the list of **Permission entries** and click **Remove**. Leave the lines for SYSTEM, Administrators, and CREATOR OWNER in the list of **Permission entries**.
- e. Click **Add...**, and type the name of the domain local group *mqm* . Click **Check Names**
- f. In response to a Windows Security window, Type the name and password of the Domain Administrator and click **OK** > **OK**.
- g. In the Permission Entry for *wmq* window, select **Full Control** in the list of **Permissions**.
- h. Click **OK** > **Apply** > **OK** > **OK** > **OK**
- i. Repeat steps e to h to add Domain Administrators.
- j. In Windows Explorer, right-click *c:\wmq* > **Share...**
- k. Click **Advanced Sharing...** and select the **Share this folder** check box. Leave the share name as *wmq* .
- l. Click **Permissions** > **Add...**, and type the name of the domain local group *mqm* ; Domain Administrators. Click **Check Names**.
- m. In response to a Windows Security window, Type the name and password of the Domain Administrator and click **OK** > **OK**.

3. Create the queue manager *QMGR* on the first domain controller, *sun* .

```
crtmqm -sax -u SYSTEM.DEAD.LETTER.QUEUE -md \\mars\wmq\data -ld \\mars\wmq\logs QMGR
```

The system response:

```
IBM MQ queue manager created.
Directory '\\mars\wmq\data\QMGR' created.
The queue manager is associated with installation 'Installation1'.
Creating or replacing default objects for queue manager 'QMGR'.
Default objects statistics : 74 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
```

4. Start the queue manager on *sun* , permitting a standby instance.

```
strmqm -x QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation 'Installation1'.
5 log records accessed on queue manager 'QMGR' during the log
replay phase.
Log replay for queue manager 'QMGR' complete.
Transaction manager state recovered for queue manager 'QMGR'.
IBM MQ queue manager 'QMGR' started using V7.1.0.0.
```

5. Create a second instance of *QMGR* on *earth* .

- a. Check the values of the Prefix and InstallationName parameters are correct for *earth* .

On *sun* , run the **dspmqinf** command:

```
dspmqinf QMGR
```

The system response:

```
QueueManager:
Name=QMGR
Directory=QMGR
Prefix=C:\ProgramData\IBM\MQ
DataPath=\\mars\wmq\data\QMGR
InstallationName=Installation1
```

- b. Copy the machine-readable form of the **QueueManager** stanza to the clipboard.

On *sun* run the **dspmqrinf** command again, with the **-o command** parameter.

```
dspmqrinf -o command QMGR
```

The system response:

```
addmqinf -s QueueManager -v Name=QMGR
-v Directory=QMGR -v Prefix="C:\ProgramData\IBM\MQ"
-v DataPath=\\mars\wmq\data\QMGR
```

- c. On *earth* run the **addmqinf** command from the clipboard to create an instance of the queue manager on *earth* .

Adjust the command, if necessary, to accommodate differences in the Prefix or InstallationName parameters.

```
addmqinf -s QueueManager -v Name= QMGR
-v Directory= QMGR -v Prefix="C:\Program Files\IBM\WebSphere MQ"
-v DataPath=\\mars\wmq\data\QMGR
```

IBM MQ configuration information added.

6. Start the standby instance of the queue manager on *earth* .

```
strmqm -x QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation 'Installation1'.
A standby instance of queue manager 'QMGR' has been started. The active
instance is running elsewhere.
```

## Results

Verify that the queue manager switches over from *sun* to *earth* :

1. On *sun* , run the command:

```
endmqm -i -r -s QMGR
```

The system response on *sun* :

```
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ended, permitting switchover to
a standby instance.
```

2. On *earth* repeatedly type the command:

```
dspmqr
```



The system responses:

```
QMNAME(QMGR) STATUS(Running as standby)
QMNAME(QMGR) STATUS(Running as standby)
QMNAME(QMGR) STATUS(Running)
```

### What to do next

To verify a multi-instance queue manager using sample programs; see “Verify the multi-instance queue manager on Windows” on page 1101.

#### Related tasks:


“Adding a second domain controller to the *wmq.example.com* domain”

Add a second domain controller to the *wmq.example.com* domain to construct a Windows domain in which to run multi-instance queue managers on domain controllers and file servers.

“Installing IBM MQ on domain controllers in the *wmq.example.com* domain” on page 1099

Put your short description here; used for first paragraph and abstract.

#### Related information:

 Windows 2000, Windows Server 2003, and Windows Server 2008 cluster nodes as domain controllers

*Adding a second domain controller to the wmq.example.com domain:*

Add a second domain controller to the *wmq.example.com* domain to construct a Windows domain in which to run multi-instance queue managers on domain controllers and file servers.

The example configuration consists of three servers:

*sun* A Windows Server 2008 server used as the first domain controller. It defines the *wmq.example.com* domain that contains *sun* , *earth* , and *mars* . It contains one instance of the multi-instance queue manager called *QMGR* .

*earth* A Windows Server 2008 used as the second domain controller IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .

*mars* A Windows Server 2008 used as the file server.

Replace the italicized names in the example, with names of your choosing.

#### Before you begin

1. Do the steps in “Creating an Active Directory and DNS domain for IBM MQ” on page 1083 to create a domain controller, *sun* , for the domain *wmq.example.com* . Change the italicized names to suit your configuration.
2. Install Windows Server 2008 on a server in the default workgroup, WORKGROUP. For the example, the server is named *earth* .

#### About this task

In this task you configure a Windows Server 2008, called *earth* , as a second domain controller in the *wmq.example.com* domain.

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

## Procedure

1. Add the domain controller, *sun.wmq.example.com* to *earth* as a DNS server.
  - a. On *earth* , log on as *earth\Administrator* and click **Start**.
  - b. Right-click **Network > Properties > Manage network connections**.
  - c. Right-click the network adapter, click **Properties**.

The system responds with the Local Area Connection Properties window listing items the connection uses.
  - d. Select the **Internet Protocol Version 4** or **Internet Protocol Version 6** from the list of items in the Local Area Connection Properties window. Click **Properties > Advanced...** and click the **DNS** tab.
  - e. Under the DNS server addresses, click **Add...**
  - f. Type the IP address of the domain controller, which is also the DNS server, and click **Add**.
  - g. Click **Append these DNS suffixes > Add...**
  - h. Type *wmq.example.com* and click **Add**.
  - i. Type *wmq.example.com* in the **DNS suffix for this connection** field.
  - j. Select **Register this connection's address in DNS** and **Use this connection's suffix in DNS registration**. Click **OK > OK > Close**
  - k. Open a command window, and type the command **ipconfig /all** to review the TCP/IP settings.
2. Log on to the domain controller, *sun* , as the local or Workgroup administrator.

If the server is already configured as a domain controller, you must log on as a domain administrator.
3. Run the Active Directory Domain Services wizard.
  - a. Click **Start > Run...** Type **dcpromo** and click **OK**.

If the Active Directory binary files are not already installed, Windows installs the files automatically.
4. Configure *earth* as the second domain controller in the *wmq.example.com* domain.
  - a. In the first window of the wizard, leave the **Use advanced mode installation** check box clear. Click **Next > Next** and click **Create Add a domain controller to an existing domain > Next**.
  - b. Type *wmq* into the **Type the name of any domain in this forest ...** field. The **Alternate credentials** radio button is clicked, click **Set...** Type in the name and password of the domain administrator and click **OK > Next > Next > Next**.
  - c. In the Additional Domain Controller Options window accept the **DNS server** and **Global catalog** options, which are selected; click **Next > Next**.
  - d. On the Directory Services Restore Mode Administrator Password, type in a **Password** and **Confirm password** and click **Next > Next**.
  - e. When prompted for **Network Credentials**, type in the password of the domain administrator. Select **Reboot on completion** in the final wizard window.
  - f. After a while, a window might open with a **DCPromo** error concerning DNS delegation; click **OK**. The server reboots.

## Results

When *earth* has rebooted, log on as Domain Administrator. Check that the *wmq.example.com* domain has been replicated to *earth* .

## What to do next

Continue with installing IBM MQ ; see "Installing IBM MQ on domain controllers in the *wmq.example.com* domain" on page 1099.

### Related tasks:

“Creating an Active Directory and DNS domain for IBM MQ” on page 1083

This task creates the domain *wmq.example.com* on a Windows 2008 domain controller called *sun* . It configures the Domain *mqm* global group in the domain, with the correct rights, and with one user.

*Installing IBM MQ on domain controllers in the wmq.example.com domain:*

Install and configure installations of IBM MQ on both domain controllers in the *wmq.example.com* domain.

Put your short description here; used for first paragraph and abstract.

The example configuration consists of three servers:

*sun* A Windows Server 2008 server used as the first domain controller. It defines the *wmq.example.com* domain that contains *sun* , *earth* , and *mars* . It contains one instance of the multi-instance queue manager called *QMGR* .

*earth* A Windows Server 2008 used as the second domain controller IBM MQ server. It contains the second instance of the multi-instance queue manager called *QMGR* .

*mars* A Windows Server 2008 used as the file server.

Replace the italicized names in the example, with names of your choosing.

### Before you begin

1. Do the steps in “Creating an Active Directory and DNS domain for IBM MQ” on page 1083 to create a domain controller, *sun* , for the domain *wmq.example.com* . Change the italicized names to suit your configuration.
2. Do the steps in “Adding a second domain controller to the *wmq.example.com* domain” on page 1097 to create a second domain controller, *earth* , for the domain *wmq.example.com* . Change the italicized names to suit your configuration.
3. See Hardware and software requirements on Windows systems for other Windows versions you can run IBM MQ on.

### About this task

Install and configure installations of IBM MQ on both domain controllers in the *wmq.example.com* domain.

### Procedure

1. Install IBM MQ on *sun* and *earth* .

For further information about running the IBM MQ for Windows installation wizard; see Installing IBM MQ server on Windows .

- a. On both *sun* and *earth* , log on as the domain administrator, *wmq\Administrator*.
- b. Run the **Setup** command on the IBM MQ for Windows installation media.  
The IBM MQ Launchpad application starts.
- c. Click **Software Requirements** to check that the prerequisite software is installed.
- d. Click **Network Configuration** > **No**.

You can configure either a domain user ID or not for this installation. The user ID that is created is a domain local user ID.

- e. Click **IBM MQ Installation**, select an installation language and click Launch IBM MQ Installer.
- f. Confirm the license agreement and click **Next** > **Next** > **Install** to accept the default configuration. Wait for the installation to complete, and click **Finish**.

If you want to change the name of the installation, install different components, configure a different directory for queue manager data and logs, or install into a different directory, click **Custom** rather than **Typical**. IBM MQ is installed, and the installer starts the “Prepare IBM MQ ” wizard.

The IBM MQ for Windows installation configures a domain local group `mqm`, and a domain group `Domain mqm`. It makes `Domain mqm` a member of `mqm`. Subsequent domain controllers in the same domain share the `mqm` and `Domain mqm` groups.

2. On both *earth* and *sun* , run the “Prepare IBM MQ ” wizard.

For further information about running the “Prepare IBM MQ ” wizard, see *Configuring IBM MQ with the Prepare IBM MQ wizard*.

- a. The IBM MQ installer runs the “Prepare IBM MQ ” automatically.

To start the wizard manually, find the shortcut to the “Prepare IBM MQ ” in the **Start > All programs > IBM MQ** folder. Select the shortcut that corresponds to the installation of IBM MQ in a multi-installation configuration.

- b. Click **Next** and leave **No** clicked in response to the question “Identify if there is a Windows 2000 or later domain controller in the network”<sup>14</sup> .
- c. In the final page of the wizard, select or clear the check boxes as you require and click **Finish**.

The “Prepare IBM MQ ” wizard creates a domain local user `MUSR_MQADMIN` on the first domain controller, and another domain local user `MUSR_MQADMIN1` on the second domain controller. The wizard creates the IBM MQ service on each controller, with `MUSR_MQADMIN` or `MUSR_MQADMIN1` as the user that logs on the service.

3. Define a user that has permission to create a queue manager.

The user must have the right to log on locally, and be a member of the domain local `mqm` group. On domain controllers, domain users do not have the right to log on locally, but administrators do. By default, no user has both these attributes. In this task, add domain administrators to the domain local `mqm` group.

- a. Open **Server Manager > Roles > Active Directory Domain Services > *wmq.example.com* > Users**.
- b. Right-click **Domain Admins > Add to a group...** and type `mqm` ; click **Check names > OK > OK**

## Results

1. Check that the “Prepare IBM MQ ” created the domain user, `MUSR_MQADMIN`:
  - a. Open **Server Manager > Roles > Active Directory Domain Services > *wmq.example.com* > Users**.
  - b. Right-click **MUSR\_MQADMIN > Properties... > Member Of**, and see that it is a member of `Domain users` and `mqm`.
2. Check that `MUSR_MQADMIN` has the right to run as a service:
  - a. Click **Start > Run...**, type the command **secpol.msc** and click **OK**.
  - b. Open **Security Settings > Local Policies > User Rights Assignments**. In the list of policies, right-click **Log on as a service > Properties** , and see `MUSR_MQADMIN` is listed as having the right to log on as a service. Click **OK**.

## What to do next

1. Do the task, “Reading and writing data and log files authorized by the local `mqm` group” on page 1108, to verify that the installation and configuration are working correctly.
2. Go back to the task, “Create a multi-instance queue manager on domain controllers” on page 1094, to complete the task of configuring a multi-instance queue manager on domain controllers.

---

14. You can configure the installation for the domain. As all users and groups on a domain controller have domain scope, it does not make any difference. It is simpler to install IBM MQ as if it is not in domain.

## Related information:

User rights required for an IBM MQ Windows Service

*Verify the multi-instance queue manager on Windows:*

Use the sample programs **amqsgbac**, **amqsphac** and **amqsmhac** to verify a multi-instance queue manager configuration. This topic provides an example configuration to verify a multi-instance queue manager configuration on Windows Server 2003.

The high availability sample programs use automatic client reconnection. When the connected queue manager fails, the client attempts to reconnect to a queue manager in the same queue manager group. The description of the samples, High availability sample programs, demonstrates client reconnection using a single instance queue manager for simplicity. You can use the same samples with multi-instance queue managers to verify a multi-instance queue manager configuration.

This example uses the multi-instance configuration described in “Create a multi-instance queue manager on domain controllers” on page 1094. Use the configuration to verify that the multi-instance queue manager switches over to the standby instance. Stop the queue manager with the **endmqm** command and use the **-s**, **switchover**, option. The client programs reconnect to the new queue manager instance and continue to work with the new instance after a slight delay.

The client is installed in a 400 MB VMware image that is running Windows 7 Service Pack 1. For security reasons, it is connected on the same VMware host-only network as the domain servers that are running the multi-instance queue manager. It is sharing the /MQHA folder, which contains the client connection table, to simplify configuration.

## Verifying failover using IBM MQ Explorer

Before using the sample applications to verify failover, run the IBM MQ Explorer on each server. Add both queue manager instances to each explorer using the **Add Remote Queue Manager > Connect directly to a multi-instance queue manager** wizard. Ensure that both instances are running, permitting standby. Close the window running the VMware image with the active instance, virtually powering off the server, or stop the active instance, allowing switchover to standby instance and reconnectable clients to reconnect.

**Note:** If you power off the server, make sure that it is not the one hosting the MQHA folder!

**Note:** The **Allow switchover to a standby instance** option might not be available on the Stop Queue Manager dialog. The option is missing because the queue manager is running as a single instance queue manager. You must have started it without the **Permit a standby instance** option. If your request to stop the queue manager is rejected, look at the Details window, possibly there is no standby instance running.

## Verifying failover using the sample programs

### Choose a server to run the active instance

You might have chosen one of the servers to host the MQHA directory or file system. If you plan to test failover by closing the VMware window running the active server, make sure that it is not the one hosting MQHA !

### On the server running the active queue manager instance

1. Modify *ipaddr1* and *ipaddr2* and save the following commands in N:\hasample.tst.

```
DEFINE QLOCAL(SOURCE) REPLACE
DEFINE QLOCAL(TARGET) REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
MCAUSER(' ') REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNAME(' ipaddr1 (1414), ipaddr2 (1414)') QMNAME(QM1) REPLACE
```

```
START CHANNEL(CHANNEL1)
DEFINE LISTENER(LISTENER.TCP) TRPTYPE(TCP) CONTROL(QMGR)
DISPLAY LISTENER(LISTENER.TCP) CONTROL
DISPLAY LSSTATUS(LISTENER.TCP) STATUS
```

**Note:** By leaving the **MCAUSER** parameter blank, the client user ID is sent to the server. The client user ID must have the correct permissions on the servers. An alternative is to set the **MCAUSER** parameter in the SVRCONN channel to the user ID you have configured on the server.

2. Open a command prompt with the path N:\ and run the command:

```
runmqsc -m QM1 < hasample.tst
```

3. Verify that the listener is running and has queue manager control, either by inspecting the output of the **runmqsc** command.

```
LISTENER(LISTENER.TCP)CONTROL(QMGR)
LISTENER(LISTENER.TCP)STATUS(RUNNING)
```

Or, using the IBM MQ Explorer that the TCPIP listener is running and has Control = Queue Manager.

### On the client

1. Map the shared directory C:\MQHA on the server to N:\ on the client.
2. Open a command prompt with the path N:\. Set the environment variable MQCHLLIB to point to the client channel definition table (CCDT) on the server:

```
SET MQCHLLIB=N:\data\QM1\@ipcc
```

3. At the command prompt type the commands:

```
start amqsgnac TARGET QM1
start amqsmnac -s SOURCE -t TARGET -m QM1
start amqspnac SOURCE QM1
```

**Note:** If you have problems, start the applications at a command prompt so that the reason code is printed out on the console, or look at the AMQERR01.LOG file in the N:\data\QM1\errors folder.

### On the server running the active queue manager instance

1. Either:
  - Close the window running the VMware image with the active server instance.
  - Using the IBM MQ Explorer, stop the active queue manager instance, allowing switchover to the standby instance and instructing re-connectable clients to reconnect.
2. The three clients eventually detect the connection is broken, and then reconnect. In this configuration, if you close the server window, it is taking about seven minutes for all three connections to be reestablished. Some connections are reestablished well before others.

### Results

```
N:\>amqspnac SOURCE QM1
Sample AMQSPHAC start
target queue is SOURCE
message <Message 1>
message <Message 2>
message <Message 3>
message <Message 4>
message <Message 5>
17:05:25 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:47 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:52 : EVENT : Connection Reconnected
message <Message 6>
message <Message 7>
message <Message 8>
message <Message 9>
```

```
N:\>amqsmhac -s SOURCE -t TARGET -m QM1
Sample AMQSMHA0 start

17:05:25 : EVENT : Connection Reconnecting (Delay: 97ms)
17:05:48 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:53 : EVENT : Connection Reconnected
```

```
N:\>amqsgnac TARGET QM1
Sample AMQSGHAC start
message <Message 1>
message <Message 2>
message <Message 3>
message <Message 4>
message <Message 5>
17:05:25 : EVENT : Connection Reconnecting (Delay: 156ms)
17:05:47 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:52 : EVENT : Connection Reconnected
message <Message 6>
message <Message 7>
message <Message 8>
message <Message 9>
```

*Secure shared queue manager data and log directories and files on Windows:*

This topic describes how you can secure a shared location for queue manager data and log files using a global alternative security group. You can share the location between different instances of a queue manager running on different servers.

Typically you do not set up a shared location for queue manager data and log files. When you install IBM MQ for Windows, the installation program creates a home directory of your choosing for any queue managers that are created on that server. It secures the directories with the local `mqm` group, and configures a user ID for the IBM MQ service to access the directories.

When you secure a shared folder with a security group, a user that is permitted to access the folder must have the credentials of the group. Suppose that a folder on a remote file server is secured with the local `mqm` group on a server called *mars*. Make the user that runs queue manager processes a member of the local `mqm` group on *mars*. The user has the credentials that match the credentials of the folder on the remote file server. Using those credentials, the queue manager is able to access its data and logs files in the folder. The user that runs queue manager processes on a different server is a member of a different local `mqm` group which does not have matching credentials. When the queue manager runs on a different server to *mars*, it cannot access the data and log files it created when it ran on *mars*. Even if you make the user a domain user, it has different credentials, because it must acquire the credentials from the local `mqm` group on *mars*, and it cannot do that from a different server.

Providing the queue manager with a global alternative security group solves the problem; see Figure 155 on page 1104. Secure a remote folder with a global group. Pass the name of the global group to the queue manager when you create it on *mars*. Pass the global group name as the alternative security group using the `-a[r]` parameter on the `crtmqm` command. If you transfer the queue manager to run on a different server, the name of the security group is transferred with it. The name is transferred in the **AccessMode** stanza in the `qm.ini` file as a `SecurityGroup`; for example:

```
AccessMode:
SecurityGroup=wmq\wmq
```

The **AccessMode** stanza in the `qm.ini` also includes the `RemoveMQMAccess`; for example:

```
AccessMode:
RemoveMQMAccess=<true\false>
```

If this attribute is specified with value true, and an access group has also been given, the local mqm group is not granted access to the queue manager data files.

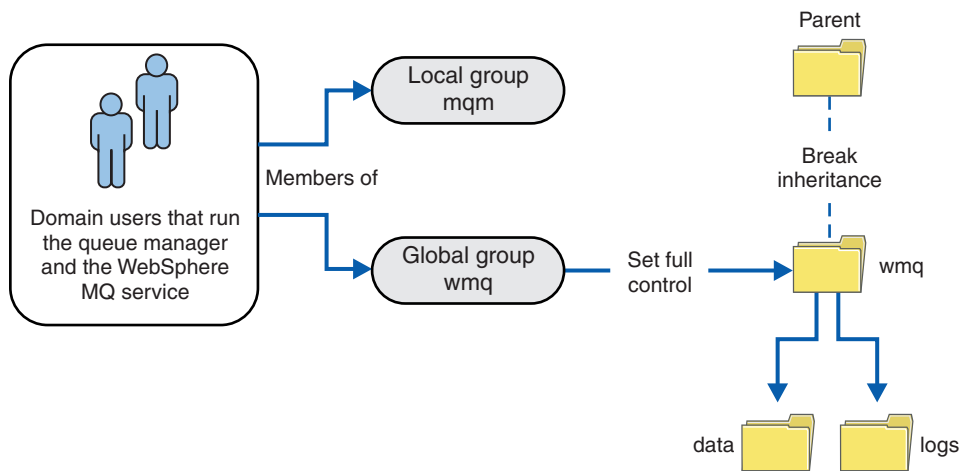


Figure 155. Securing queue manager data and logs using an alternative global security group (1)

For the user ID that queue manager processes are to run with to have the matching credentials of the global security group, the user ID must also have global scope. You cannot make a local group or principal a member of a global group. In Figure 155, the users that run the queue manager processes are shown as domain users.

If you are deploying many IBM MQ servers, the grouping of users in Figure 155 is not convenient. You would need to repeat the process of adding users to local groups for every IBM MQ server. Instead, create a Domain mqm global group on the domain controller, and make the users that run IBM MQ members of the Domain mqm group; see Figure 156 on page 1105. When you install IBM MQ as a domain installation, the "Prepare IBM MQ" wizard automatically makes the Domain mqm group a member of the local mqm group. The same users are in both the global groups Domain mqm and wmq.

**Tip:** The same users can run IBM MQ on different servers, but on an individual server you must have different users to run IBM MQ as a service, and run interactively. You must also have different users for every installation on a server. Typically, therefore Domain mqm contains a number of users.



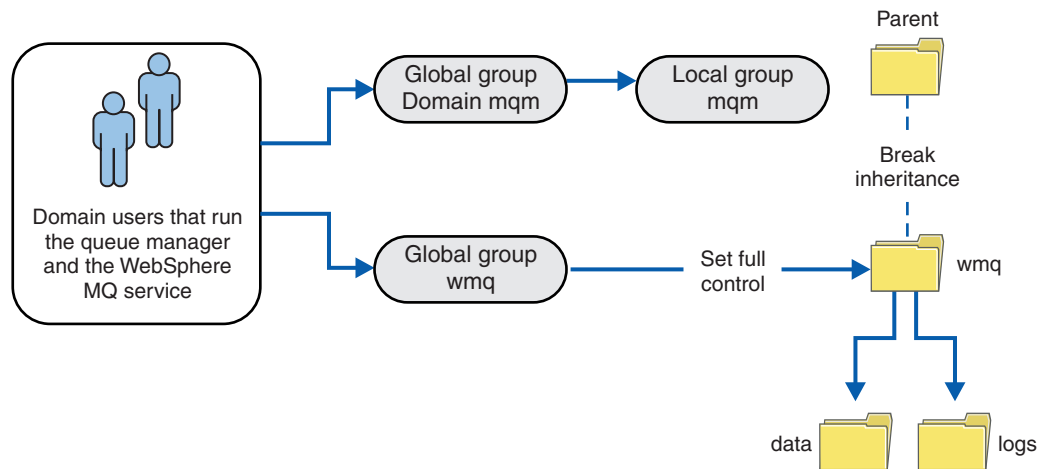


Figure 156. Securing queue manager data and logs using an alternative global security group (2)

The organization in Figure 156 is unnecessarily complicated as it stands. The arrangement has two global groups with identical members. You might simplify the organization, and define only one global group; see Figure 157.

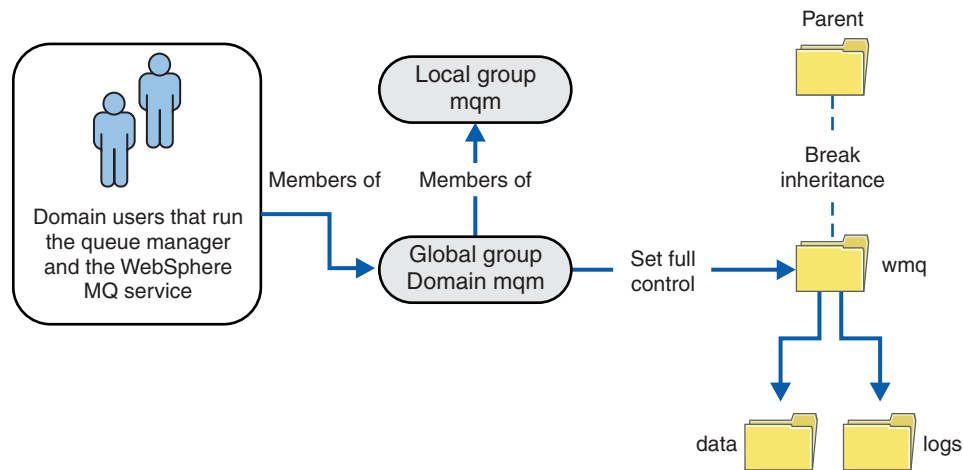


Figure 157. Securing queue manager data and logs using an alternative global security group (3)

Alternatively, you might need a finer degree of access control, with different queue managers restricted to being able to access different folders; see Figure 158 on page 1106. In Figure 158 on page 1106, two groups of domain users are defined, in separate global groups to secure different queue manager log and data files. Two different local mqm groups are shown, which must be on different IBM MQ servers. In this example, the queue managers are partitioned into two sets, with different users allocated to the two sets. The two sets might be test and production queue managers. The alternate security groups are called wmq1 and wmq2. You must manually add the global groups wmq1 and wmq2 to the correct queue managers according to whether they are in the test or production department. The configuration cannot take advantage that the installation of IBM MQ propagates Domain mqm to the local mqm group as in Figure 157, because there are two groups of users.

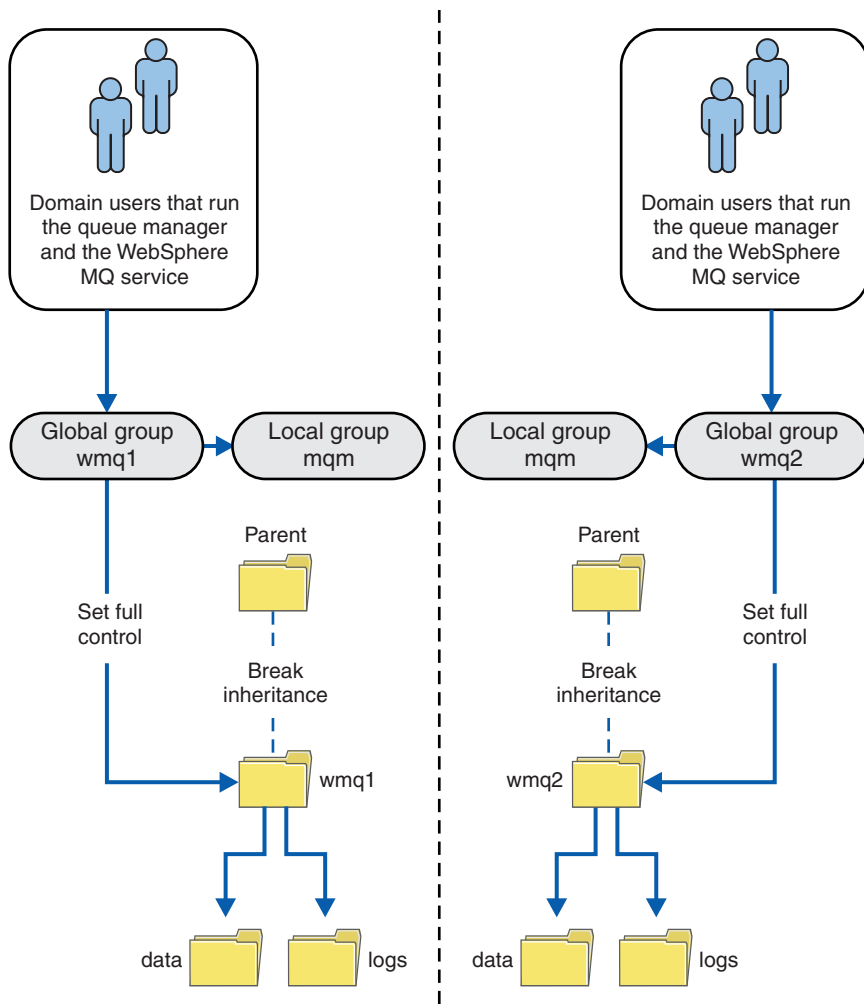


Figure 158. Securing queue manager data and logs using an alternative global security principal (4)

An alternative way to partition two departments would be to place them in two windows domains. In that case, you could return to using the simpler model shown in Figure 157 on page 1105.

*Secure unshared queue manager data and log directories and files on Windows:*

This topic describes how you can secure an alternative location for queue manager data and log files, both by using the local mqm group and an alternative security group.

Typically you do not set up an alternative location for queue manager data and log files. When you install IBM MQ for Windows, the installation program creates a home directory of your choosing for any queue managers that are created. It secures the directories with the local mqm group, and configures a user ID for the IBM MQ service to access the directories.

Two examples demonstrate how to configure access control for IBM MQ. The examples show how to create a queue manager with its data and logs in directories that are not on the data and log paths created by the installation. In the first example, "Reading and writing data and log files authorized by the local mqm group" on page 1108, you permit access to the queue and log directories by authorizing by the local mqm group. The second example, "Reading and writing data and log files authorized by an alternative local security group" on page 1111, differs in that access to the directories is authorized by an alternative security group. When the directories are accessed by a queue manager running on only one

server, securing the data and log files with the alternative security group gives you the choice of securing different queue managers with different local groups or principals. When the directories are accessed by a queue manager running on different servers, such as with a multi-instance queue manager, securing the data and log files with the alternate security group is the only choice; see “Secure shared queue manager data and log directories and files on Windows” on page 1103.

Configuring the security permissions of queue manager data and log files is not a common task on Windows. When you install IBM MQ for Windows, you either specify directories for queue manager data and logs, or accept the default directories. The installation program automatically secures these directories with the local `mqm` group, giving it full control permission. The installation process makes sure the user ID that runs queue manager is a member of the local `mqm` group. You can modify the other access permissions on the directories to meet your access requirements.

If you move the data and log files directory to new locations, you must configure the security of the new locations. You might change the location of the directories if you back up a queue manager and restore it onto a different computer, or if you change the queue manager to be a multi-instance queue manager. You have a choice of two ways of securing the queue manager data and log directories in their new location. You can secure the directories by restricting access to the local `mqm` group, or you can restrict access to any security group of your choosing.

It takes the least number of steps to secure the directories using the local `mqm` group. Set the permissions on the data and log directories to allow the local `mqm` group full control. A typical approach is to copy the existing set of permissions, removing inheritance from the parent. You can then remove or restrict the permissions of other principals.

If you run the queue manager under a different user ID to the service set up by the Prepare IBM MQ wizard, that user ID must be a member of the local `mqm` group. The task, “Reading and writing data and log files authorized by the local `mqm` group” on page 1108, takes you through the steps.

You can also secure queue manager data and log files using an alternative security group. The process of securing the queue manager data and log files with the alternative security group has a number of steps that refer to Figure 159. The local group, `wmq`, is an example of an alternative security group.

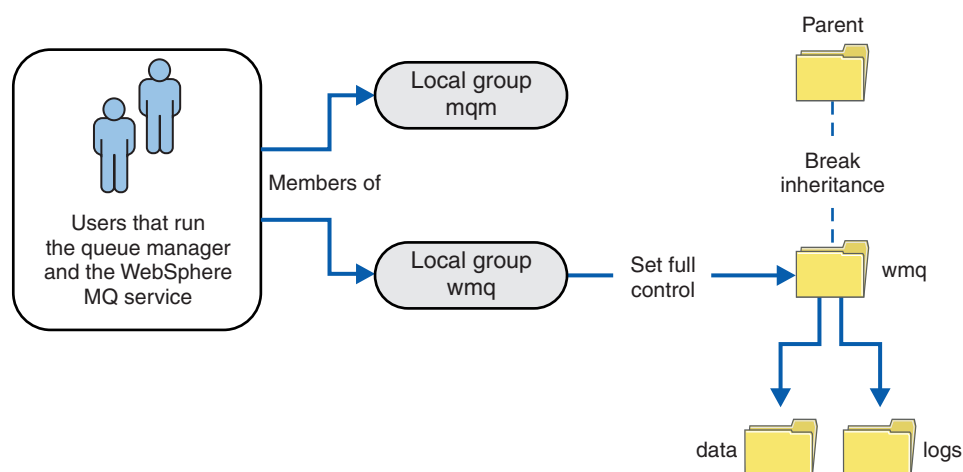


Figure 159. Securing queue manager data and logs using an alternative local security group, `wmq`

1. Either create separate directories for the queue manager data and logs, a common directory, or a common parent directory.

2. Copy the existing set of inherited permissions for the directories, or parent directory, and modify them according to your needs.
3. Secure the directories that are to contain the queue manager and logs by giving the alternative group, `wmq`, full control permission to the directories.
4. Give all user IDs that run queue manager processes the credentials of the alternative security group or principal:
  - a. If you define a user as the alternative security principal, the user must be same user as the queue manager is going to run under. The user must be a member of the local `mqm` group.
  - b. If you define a local group as the alternative security group, add the user that the queue manager is going to run under to the alternative group. The user must also be a member of the local `mqm` group.
  - c. If you define an global group as the alternative security group, then see “Secure shared queue manager data and log directories and files on Windows” on page 1103.
5. Create the queue manager specifying the alternative security group or principal on the `crtmqm` command, with the `-a` parameter.

*Reading and writing data and log files authorized by the local mqm group:*

The task illustrates how to create a queue manager with its data and logs files stored in any directory of your choosing. Access to the files is secured by the local `mqm` group. The directory is not shared.

### **Before you begin**

1. Install IBM MQ for Windows as the primary installation.
2. Run the “Prepare IBM MQ ” wizard. For this task, configure the installation either to run with a local user ID, or a domain user ID. Eventually, to complete all the tasks in “Windows domains and multi-instance queue managers” on page 1079, the installation must be configured for a domain.
3. Log on with Administrator authority to perform the first part of the task.

### **About this task**

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

On Windows, you can create the default data and log paths for an IBM MQ for Windows in any directories of your choosing. The installation and configuration wizard automatically gives the local `mqm` group, and the user ID that is running the queue manager processes, access to the directories. If you create a queue manager specifying different directories for queue manager data and log files, you must configure full control permission to the directories.

In this example, you give the queue manager full control over its data and log files by giving the local `mqm` group permission to the directory `c:\wmq` .

The `crtmqm` command creates a queue manager that starts automatically when the workstation starts using the IBM MQ service.

The task is illustrative; it uses specific values that you can change. The values you can change are in italics. At the end of the task, follow the instructions to remove all the changes you made.

### **Procedure**

1. Open a command prompt.
2. Type the command:

```
md c:\wmq\data, c:\wmq\logs
```

3. Set the permissions on the directories to allow the local mqm group read and write access.

```
cacls c:\wmq/T /E /G mqm:F
```

The system response:

```
processed dir: c:\wmq
processed dir: c:\wmq\data
processed dir: c:\wmq\logs
```

4. Optional: Switch to a user ID that is a member of the local mqm group.

You can continue as Administrator, but for a realistic production configuration, continue with a user ID with more restricted rights. The user ID must at least be a member of the local mqm group.

If the IBM MQ installation is configured as part of a domain, make the user ID a member of the Domain mqm group. The "Prepare IBM MQ " wizard makes the Domain mqm global group a member of the local mqm group, so you do not have to make the user ID directly a member of the local mqm group.

5. Create the queue manager.

```
crtmqm -sax -u SYSTEM.DEAD.LETTER.QUEUE -md c:\wmq\data -ld c:\wmq\logs QMGR
```

The system response:

```
IBM MQ queue manager created.
Directory 'c:\wmq\data\QMGR' created.
The queue manager is associated with installation '1'
Creating or replacing default objects for queue manager 'QMGR'
Default objects statistics : 74 created. 0 replaced.
Completing setup.
Setup completed.
```

6. Check that the directories created by the queue manager are in the c:\wmq directory.

```
dir c:\wmq/D /B /S
```

7. Check that the files have read and write, or full control permission for the local mqm group.

```
cacls c:\wmq*.*
```

## What to do next

Test the queue manager by putting and getting a message to a queue.

1. Start the queue manager.

```
strmqm QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation '1'.
5 log records accessed on queue manager 'QMGR' during the log
replay phase.
Log replay for queue manager 'QMGR' complete.
Transaction manager state recovered for queue manager 'QMGR'.
IBM MQ queue manager 'QMGR' started using V7.1.0.0.
```

2. Create a test queue.

```
echo define qlocal(QTEST) | runmqsc QMGR
```

The system response:

5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.  
Starting MQSC for queue manager QMGR.

```
1 : define qlocal(QTEST)
AMQ8006: IBM MQ queue created.
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

3. Put a test message using the sample program **amqsput**.

```
echo 'A test message' | amqsput QTEST QMGR
```

The system response:

```
Sample AMQSPUT0 start
target queue is QTEST
Sample AMQSPUT0 end
```

4. Get the test message using the sample program **amqsget**.

```
amqsget QTEST QMGR
```

The system response:

```
Sample AMQSGET0 start
message <A test message>
Wait 15 seconds ...
no more messages
Sample AMQSGET0 end
```

5. Stop the queue manager.

```
endmqm -i QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ended.
```

6. Delete the queue manager.

```
dltmqm QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' deleted.
```

7. Delete the directories you created.

**Tip:** Add the /Q option to the commands to prevent the command prompting to delete each file or directory.

```
del /F /S C:\wmq*.*
rmdir /S C:\wmq
```

### Related concepts:

“Windows domains and multi-instance queue managers” on page 1079

A multi-instance queue manager on Windows requires its data and logs to be shared. The share must be accessible to all instances of the queue manager running on different servers or workstations. Configure the queue managers and share as part of a Windows domain. The queue manager can run on a domain workstation or server, or on the domain controller.

### Related tasks:

“Reading and writing shared data and log files authorized by an alternative global security group” on page 1091

This task shows how to use the **-a** flag on the **crtmqm** command. The **-a** flag gives the queue manager access to its log and data files on a remote file share using the alternative security group.

“Create a multi-instance queue manager on domain workstations or servers” on page 1080

An example shows how to set up a multi-instance queue manager on Windows on a workstation or a server that is part of a Windows domain. The server does not have to be a domain controller. The setup demonstrates the concepts involved, rather than being production scale. The example is based on Windows Server 2008. The steps might differ on other versions of Windows Server.

*Reading and writing data and log files authorized by an alternative local security group:*

This task shows how to use the **-a** flag on the **crtmqm** command. The flag provides the queue manager with an alternative local security group to give it access to its log and data files.

### Before you begin

1. Install IBM MQ for Windows as the primary installation.
2. Run the “Prepare IBM MQ ” wizard. For this task, configure the installation either to run with a local user ID, or a domain user ID. Eventually, to complete all the tasks in “Windows domains and multi-instance queue managers” on page 1079, the installation must be configured for a domain.
3. Log on with Administrator authority to perform the first part of the task.

### About this task

This task is one of a set of related tasks that illustrate accessing queue manager data and log files. The tasks show how to create a queue manager authorized to read and write data and log files that are stored in a directory of your choosing. They accompany the task, “Windows domains and multi-instance queue managers” on page 1079.

On Windows, you can create the default data and log paths for an IBM MQ for Windows in any directories of your choosing. The installation and configuration wizard automatically gives the local **mqm** group, and the user ID that is running the queue manager processes, access to the directories. If you create a queue manager specifying different directories for queue manager data and log files, you must configure full control permission to the directories.

In this example, you provide the queue manager with an alternative security local group that has full control authorization to the directories. The alternative security group gives the queue manager permission to manage files in the directory. The primary purpose of the alternate security group is to authorize an alternate security global group. Use an alternate security global group to set up a multi-instance queue manager. In this example, you configure a local group to familiarize yourself with the use of an alternate security group without installing IBM MQ in a domain. It is unusual to configure a local group as an alternative security group.

The **crtmqm** command creates a queue manager that starts automatically when the workstation starts using the IBM MQ service.

The task is illustrative; it uses specific values that you can change. The values you can change are in italics. At the end of the task, follow the instructions to remove all the changes you made.

## Procedure

### 1. Set up an alternative security group.

The alternative security group is typically a domain group. In the example, you create a queue manager that uses a local alternate security group. With a local alternate security group, you can do the task with an IBM MQ installation that is not part of a domain.

- a. Run the **lusrmgr.msc** command to open the Local Users and Groups window.
- b. Right-click **Groups > New Group...**
- c. In the **Group name** field, type *altnmqm* and click **Create > Close**.
- d. Identify the user ID that runs the IBM MQ service.
  - 1) Click **Start > Run...**, type *services.msc* and click **OK**.
  - 2) Click the IBM MQ service in the list of services, and click the Log On tab.
  - 3) Remember the user ID and close the Services Explorer.
- e. Add the user ID that runs the IBM MQ service to the *altnmqm* group. Also add the user ID that you log on with to create a queue manager, and run it interactively.

Windows checks the authority of the queue manager to access the data and logs directories by checking the authority of the user ID that is running queue manager processes. The user ID must be a member, directly or indirectly through a global group, of the *altnmqm* group that authorized the directories.

If you installed IBM MQ as part of a domain, and are going to do the tasks in “Create a multi-instance queue manager on domain workstations or servers” on page 1080, the domain user IDs created in “Creating an Active Directory and DNS domain for IBM MQ” on page 1083 are *wmquser1* and *wmquser2* .

If you did not install the queue manager as part of a domain, the default local user ID that runs the IBM MQ service is *MUSR\_MQADMIN*. If you intend to do the tasks without Administrator authority, create a user that is a member of the local *mqm* group.

Follow these steps to add *wmquser1* and *wmquser2* to *altnmqm* . If your configuration is different, substitute your names for the user IDs and group.

- 1) In the list of groups right-click **altnmqm > Properties > Add...**
- 2) In the Select Users, Computers, or Groups window type *wmquser1 ; wmquser2* and click **Check Names**.
- 3) Type the name and password of a domain administrator in the Windows Security window, then click **OK > OK > Apply > OK**.

### 2. Open a command prompt.

### 3. Restart the IBM MQ service.

You must restart the service so that the user ID it runs under acquires the additional security credentials you configured for it.

Type the commands:

```
endmqsvc
strmqsvc
```

The system responses:

```
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
The MQ service for installation 'Installation1' ended successfully.
```

And:



5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.  
The MQ service for installation 'Installation1' started successfully.

4. Type the command:

```
md c:\wmq\data, c:\wmq\logs
```

5. Set the permissions on the directories to allow the local user *user* read and write access.

```
cacls c:\wmq/T /E /G altmqm:F
```

The system response:

```
processed dir: c:\wmq
processed dir: c:\wmq\data
processed dir: c:\wmq\logs
```

6. Optional: Switch to a user ID that is a member of the local *mqm* group.

You can continue as Administrator, but for a realistic production configuration, continue with a user ID with more restricted rights. The user ID must at least be a member of the local *mqm* group.

If the IBM MQ installation is configured as part of a domain, make the user ID a member of the Domain *mqm* group. The "Prepare IBM MQ " wizard makes the Domain *mqm* global group a member of the local *mqm* group, so you do not have to make the user ID directly a member of the local *mqm* group.

7. Create the queue manager.

```
crtmqm -a altmqm -sax -u SYSTEM.DEAD.LETTER.QUEUE -md c:\wmq\data -ld c:\wmq\logs QMGR
```

The system response:

```
IBM MQ queue manager created.
Directory 'c:\wmq\data\QMGR' created.
The queue manager is associated with installation '1'
Creating or replacing default objects for queue manager 'QMGR'
Default objects statistics : 74 created. 0 replaced.
Completing setup.
Setup completed.
```

8. Check that the directories created by the queue manager are in the *c:\wmq* directory.

```
dir c:\wmq/D /B /S
```

9. Check that the files have read and write, or full control permission for the local *mqm* group.

```
cacls c:\wmq*.*
```

## What to do next

Test the queue manager by putting and getting a message to a queue.

1. Start the queue manager.

```
strmqm QMGR
```

The system response:

```
IBM MQ queue manager 'QMGR' starting.
The queue manager is associated with installation '1'.
5 log records accessed on queue manager 'QMGR' during the log
replay phase.
Log replay for queue manager 'QMGR' complete.
Transaction manager state recovered for queue manager 'QMGR'.
IBM MQ queue manager 'QMGR' started using V7.1.0.0.
```

2. Create a test queue.

```
echo define qlocal(QTEST) | runmqsc QMGR
```

The system response:

```
5724-H72 (C) Copyright IBM Corp. 1994, 2011. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QMGR.
```

```
1 : define qlocal(QTEST)
AMQ8006: IBM MQ queue created.
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

3. Put a test message using the sample program **amqsput**.  
echo 'A test message' | amqsput QTEST QMGR

The system response:

```
Sample AMQSPUT0 start
target queue is QTEST
Sample AMQSPUT0 end
```

4. Get the test message using the sample program **amqsget**.  
amqsget QTEST QMGR

The system response:

```
Sample AMQSGET0 start
message <A test message>
Wait 15 seconds ...
no more messages
Sample AMQSGET0 end
```

5. Stop the queue manager.  
endmqm -i QMGR

The system response:

```
IBM MQ queue manager 'QMGR' ending.
IBM MQ queue manager 'QMGR' ended.
```

6. Delete the queue manager.  
dltmqm QMGR

The system response:

```
IBM MQ queue manager 'QMGR' deleted.
```

7. Delete the directories you created.

**Tip:** Add the /Q option to the commands to prevent the command prompting to delete each file or directory.

```
del /F /S C:\wmq*.*
rmdir /S C:\wmq
```

Create a multi-instance queue manager on Linux:

An example shows how to set up a multi-instance queue manager on Linux. The setup is small to illustrate the concepts involved. The example is based on Linux Red Hat Enterprise 5. The steps differ on other UNIX platforms.

### About this task

The example is set up on a 2 GHz notebook computer with 3 GB RAM running Windows 7 Service Pack 1. Two VMware virtual machines, Server1 and Server2, run Linux Red Hat Enterprise 5 in 640 MB images. Server1 hosts the network file system (NFS), the queue manager logs and an HA instance. It is not usual practice for the NFS server also to host one of the queue manager instances; this is to simplify the example. Server2 mounts Server1's queue manager logs with a standby instance. A WebSphere MQ MQI client is installed on an additional 400 MB VMware image that runs Windows 7 Service Pack 1 and runs the sample high availability applications. All the virtual machines are configured as part of a VMware host-only network for security reasons.

**Note:** You should put only queue manager data on an NFS server. On the NFS, use the following three options with the mount command to make the system secure:

#### noexec

By using this option, you stop binary files from being run on the NFS, which prevents a remote user from running unwanted code on the system.

#### nosuid

By using this option, you prevent the use of the set-user-identifier and set-group-identifier bits, which prevents a remote user from gaining higher privileges.

**nodev** By using this option, you stop character and block special devices from being used or defined, which prevents a remote user from getting out of a chroot jail.

### Procedure

1. Log in as root.
2. Follow the instructions in Installing IBM MQ to install IBM MQ, create the mqm user and group, and define /var/mqm.
3. Complete the task Verifying shared file system behavior to check that the file system supports multi-instance queue managers.
4. For Server1, complete the following step:
  - a. Create log and data directories in a common folder, /MQHA, that is to be shared. For example:
    - 1) **mkdir /MQHA**
    - 2) **mkdir /MQHA/logs**
    - 3) **mkdir /MQHA/qmgrs**
5. For Server2, complete the following step:
  - a. Create the folder, /MQHA, to mount the shared file system. Keep the path the same as on Server1. For example:
    - 1) **mkdir /MQHA**
6. Ensure that the MQHA directories are owned by user and group mqm, and the access permissions are set to rwx for user and group. For example **ls -al** displays `drwxrwxr-x mqm mqm 4096 Nov 27 14:38 MQDATA .`
  - a. **chown -R mqm:mqm /MQHA**
  - b. **chmod -R ug+rwx /MQHA**
7. Create the queue manager by entering the following command: **crtmqm -ld /MQHA/logs -md /MQHA/qmgrs QM1**

8. Add<sup>15</sup> /MQHA \*(rw, sync, no\_wdelay, fsid=0) to /etc/exports
9. For Server1, complete the following steps:
  - a. Start the NFS daemon: `/etc/init.d/ nfs start`
  - b. Copy the queue manager configuration details from Server1:
 

```
dspmqlnf -o command QM1
```

and copy the result to the clipboard:

```
addmqinf -s QueueManager
-v Name=QM1
-v Directory=QM1
-v Prefix=/var/mqm
-v DataPath=/MQHA/qmgrs/QM1
```
10. For Server2, complete the following steps:
  - a. Mount the exported file system /MQHA by entering the following command: `mount -t nfs4 -o hard, intr Server1:/ /MQHA`
  - b. Paste the queue manager configuration command into Server2:
 

```
addmqinf -s QueueManager
-v Name=QM1
-v Directory=QM1
-v Prefix=/var/mqm
-v DataPath=/MQHA/qmgrs/QM1
```
11. Start the queue manager instances, in either order, with the `-x` parameter: `strmqm -x QM1`.  
The command used to start the queue manager instances must be issued from the same IBM MQ installation as the `addmqinf` command. To start and stop the queue manager from a different installation, you must first set the installation associated with the queue manager using the `setmqm` command. For more information, see `setmqm`.

*Verifying the multi-instance queue manager on Linux:*

Use the sample programs `amqsgshac`, `amqspshac` and `amqsmshac` to verify a multi-instance queue manager configuration. This topic provides an example configuration to verify a multi-instance queue manager configuration on Linux Red Hat Enterprise 5.

The high availability sample programs use automatic client reconnection. When the connected queue manager fails, the client attempts to reconnect to a queue manager in the same queue manager group. The description of the samples, High availability sample programs, demonstrates client reconnection using a single instance queue manager for simplicity. You can use the same samples with multi-instance queue managers to verify a multi-instance queue manager configuration.

The example uses the multi-instance configuration described in "Create a multi-instance queue manager on Linux" on page 1115. Use the configuration to verify that the multi-instance queue manager switches over to the standby instance. Stop the queue manager with the `endmqm` command and use the `-s`, switchover, option. The client programs reconnect to the new queue manager instance and continue to work with the new instance after a slight delay.

In the example, the client is running on a Windows 7 Service Pack 1 system. The system is hosting two VMware Linux servers that are running the multi-instance queue manager.

### Verifying failover using IBM MQ Explorer

Before using the sample applications to verify failover, run the IBM MQ Explorer on each server. Add both queue manager instances to each explorer using the **Add Remote Queue Manager > Connect**

---

15. The '\*' allows all machines that can reach this one mount /MQHA for read/write. Restrict access on a production machine.

directly to a multi-instance queue manager wizard. Ensure that both instances are running, permitting standby. Close the window running the VMware image with the active instance, virtually powering off the server, or stop the active instance, allowing switchover to standby instance.

**Note:** If you power off the server, make sure that it is not the one hosting /MQHA !

**Note:** The **Allow switchover to a standby instance** option might not be available on the Stop Queue Manager dialog. The option is missing because the queue manager is running as a single instance queue manager. You must have started it without the **Permit a standby instance** option. If your request to stop the queue manager is rejected, look at the Details window, it is possibly because there is no standby instance running.

### Verifying failover using the sample programs

#### Choose a server to be to run the active instance

You might have chosen one of the servers to host the MQHA directory or file system. If you plan to test failover by closing the VMware window running the active server, make sure that it is not the one hosting MQHA !

#### On the server running the active queue manager instance

**Note:** Running the SVRCONN channel with the MCAUSER set to mqm, is a convenience to reduce the number of configuration steps in the example. If another user ID is chosen, and your system is set up differently to the one used in the example, you might experience access permission problems. Do not use mqm as a MCAUSER on an exposed system; it is likely to compromise security greatly.

1. Modify *ipaddr1* and *ipaddr2* and save the following commands in /MQHA/hasamples.tst.

```
DEFINE QLOCAL(SOURCE) REPLACE
DEFINE QLOCAL(TARGET) REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
MCAUSER('mqm') REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(' ipaddr1 (1414), ipaddr2
(1414)') QMNAME(QM1) REPLACE
START CHANNEL(CHANNEL1)
DEFINE LISTENER(LISTENER.TCP) TRPTYPE(TCP) CONTROL(QMGR)
DISPLAY LISTENER(LISTENER.TCP) CONTROL
START LISTENER(LISTENER.TCP)
DISPLAY LSSTATUS(LISTENER.TCP) STATUS
```

2. Open a terminal window with the path /MQHA and run the command:

```
runmqsc -m QM1 < hasamples.tst
```

3. Verify that the listener is running and has queue manager control, either by inspecting the output of the **runmqsc** command.

```
LISTENER(LISTENER.TCP)CONTROL(QMGR)
LISTENER(LISTENER.TCP)STATUS(RUNNING)
```

Or, using the IBM MQ Explorer that the TCPIP listener is running and has Control = Queue Manager.

#### On the client

1. Copy the client connection table AMQCLCHL.TAB from /MQHA/qmgrs/QM1.000/@ipcc on the server to C:\ on the client.
2. Open a command prompt with the path C:\ and set the environment variable MQCHLLIB to point to the client channel definition table (CCDT)  
SET MQCHLLIB=C:\
3. At the command prompt type the commands:

```
start amqsgnac TARGET QM1
start amqsmnac -s SOURCE -t TARGET -m QM1
start amqspnac SOURCE QM1
```

### On the server running the active queue manager instance

1. Either:
  - Close the window running the VMware image with the active server instance.
  - Using the IBM MQ Explorer, stop the active queue manager instance, allowing switchover to the standby instance and instructing reconnectable clients to reconnect.
2. The three clients eventually detect the connection is broken, and then reconnect. In this configuration, if you close the server window, it is taking about seven minutes for all three connections to be reestablished. Some connections are reestablished well before others.

### Results

```
N:\>amqspnac SOURCE QM1
Sample AMQSPNAC start
target queue is SOURCE
message <Message 1>
message <Message 2>
message <Message 3>
message <Message 4>
message <Message 5>
17:05:25 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:47 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:52 : EVENT : Connection Reconnected
message <Message 6>
message <Message 7>
message <Message 8>
message <Message 9>
```

```
N:\>amqsmnac -s SOURCE -t TARGET -m QM1
Sample AMQSMNAC start

17:05:25 : EVENT : Connection Reconnecting (Delay: 97ms)
17:05:48 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:53 : EVENT : Connection Reconnected
```

```
N:\>amqsgnac TARGET QM1
Sample AMQSGNAC start
message <Message 1>
message <Message 2>
message <Message 3>
message <Message 4>
message <Message 5>
17:05:25 : EVENT : Connection Reconnecting (Delay: 156ms)
17:05:47 : EVENT : Connection Reconnecting (Delay: 0ms)
17:05:52 : EVENT : Connection Reconnected
message <Message 6>
message <Message 7>
message <Message 8>
message <Message 9>
```

## Deleting a multi-instance queue manager:

To delete a multi-instance queue manager completely, you need to use the **dltmqm** command to delete the queue manager, and then remove instances from other servers using either the **rmvmqinf** or **dltmqm** commands.

Run the **dltmqm** command to delete a queue manager that has instances defined on other servers, on any server where that queue manager is defined. You do not need to run the **dltmqm** command on the same server that you created it on. Then run the **rmvmqinf** or **dltmqm** command on all the other servers which have a definition of the queue manager.

You can only delete a queue manager when it is stopped. At the time you delete it no instances are running, and the queue manager, strictly speaking, is neither a single or a multi-instance queue manager; it is simply a queue manager that has its queue manager data and logs on a remote share. When you delete a queue manager, its queue manager data and logs are deleted, and the queue manager stanza is removed from the `mq5.ini` file on the server on which you issued the **dltmqm** command. You need to have access to the network share containing the queue manager data and logs when you delete the queue manager.

On other servers where you have previously created instances of the queue manager there are also entries in the `mq5.ini` files on those servers. You need to visit each server in turn, and remove the queue manager stanza by running the command **rmvmqinf** *Queue manager stanza name*.

On UNIX and Linux systems, if you have placed a common `mq5.ini` file in network storage and referenced it from all the servers by setting the `AMQ_MQS_INI_LOCATION` environment variable on each server, then you need to delete the queue manager from only one of its servers as there is only one `mq5.ini` file to update.

### Example

#### First server

```
dltmqm QM1
```

#### Other servers where instances are defined

```
rmvmqinf QM1 , or
```

```
dltmqm QM1
```

## Starting and stopping a multi-instance queue manager:

Starting and stopping a queue manager configured either as a single instance or a multi-instance queue manager.

When you have defined a multi-instance queue manager on a pair of servers, you can run the queue manager on either server, either as a single instance queue manager, or as a multi-instance queue manager.

To run a multi-instance queue manager, start the queue manager on one of the servers using the **strmqm -x QM1** command; the `-x` option permits the instance to failover. It becomes the *active instance*. Start the standby instance on the other server using the same **strmqm -x QM1** command; the `-x` option permits the instance to start as a standby.

The queue manager is now running with one active instance that is processing all requests, and a standby instance that is ready to take over if the active instance fails. The active instance is granted exclusive access to the queue manager data and logs. The standby waits to be granted exclusive access to the queue manager data and logs. When the standby is granted exclusive access, it becomes the active instance.

You can also manually switch control to the standby instance by issuing the **endmqm -s** command on the active instance. The **endmqm -s** command shuts down the active instance without shutting down the standby. The exclusive access lock on the queue manager data and logs is released, and the standby takes over.

You can also start and stop a queue manager configured with multiple instances on different servers as a single instance queue manager. If you start the queue manager without using the **-x** option on the **strmqm** command, the instances of the queue manager configured on other machines are prevented from starting as standby instances. If you attempt to start another instance you receive the response that the queue manager instance is not permitted to run as a standby.

If you stop the active instance of a multi-instance queue manager using the **endmqm** command without the **-s** option, then the active and standby instances both stop. If you stop the standby instance using the **endmqm** command with the **-x** option, then it stops being a standby, and the active instance continues running. You cannot issue **endmqm** without the **-x** option on the standby.

Only two queue manager instances can run at the same time; one is the active instance, and the other is a standby instance. If you start two instances at the same time, IBM MQ has no control over which instance becomes the active instance; it is determined by the network file system. The first instance to acquire exclusive access to the queue manager data becomes the active instance.

**Note:** Before you restart a failed queue manager, you must disconnect your applications from that instance of the queue manager. If you do not, the queue manager might not restart correctly.

#### **Shared file system:**

A multi-instance queue manager uses a networked file system to manage queue manager instances.

A multi-instance queue manager automates failover using a combination of file system locks and shared queue manager data and logs. Only one instance of a queue manager can have exclusive access to the shared queue manager data and logs. When it gets access it becomes the active instance. The other instance that does not succeed in getting exclusive access waits as a standby instance until the queue manager data and logs become available.

The networked file system is responsible for releasing the locks it holds for the active queue manager instance. If the active instance fails in some way, the networked file system releases the locks it is holding for the active instance. As soon as the exclusive lock is released, a standby queue manager waiting for the lock attempts to acquire it. If it succeeds, it becomes the active instance and has exclusive access to the queue manager data and logs on the shared file system. It then continues to start.

The related topic, Planning file system support describes how to set up and check that your file system supports multi-instance queue managers.

A multi-instance queue manager does not protect you against a failure in the file system. There are a number of ways to protect your data.

- Invest in reliable storage, such as redundant disk arrays (RAID), and include them in a networked file system that has network resilience.
- Back up IBM MQ linear logs to alternative media, and if your primary log media fails, recover using the logs on the alternative media. You can use a backup queue manager to administer this process.



## Multiple queue manager instances:

A multi-instance queue manager is resilient because it uses a standby queue manager instance to restore queue manager availability after failure.

Replicating queue manager instances is a very effective way to improve the availability of queue manager processes. Using a simple availability model, purely for illustration: if the reliability of one instance of a queue manager is 99% (over one year, cumulative downtime is 3.65 days) then adding another instance of the queue manager increases the availability to 99.99% (over one year, cumulative downtime of about an hour).

This is too simple a model to give you practical numeric estimates of availability. To model availability realistically, you need to collect statistics for the mean time between failures (MTBF) and the mean time to repair (MTTR), and the probability distribution of time between failures and of repair times.

The term, multi-instance queue manager, refers to the combination of active and standby instances of the queue manager that share the queue manager data and logs. Multi-instance queue managers protect you against the failure of queue manager processes by having one instance of the queue manager active on one server, and another instance of the queue manager on standby on another server, ready to take over automatically should the active instance fail.

### Failover or switchover:

A standby queue manager instance takes over from the active instance either on request (switchover), or when the active instance fails (failover).

- *Switchover* takes place when a standby instance starts in response to the **endmqm -s** command being issued to the active queue manager instance. You can specify the **endmqm** parameters **-c**, **-i** or **-p** to control how abruptly the queue manager is stopped.

**Note:** Switchover only takes place if a standby queue manager instance is already started. The **endmqm -s** command releases the active queue manager lock and permits switchover: it does not start a standby queue manager instance.

- *Failover* occurs when the lock on queue manager data held by the active instance is released because the instance appeared to stop unexpectedly (that is, without an **endmqm** command being issued).

When the standby instance takes over as the active instance, it writes a message to the queue manager error log.

Reconnectable clients are automatically reconnected when a queue manager fails or switches over. You do not need to include the **-r** flag on the **endmqm** command to request client reconnection. Automatic client reconnect is not supported by IBM MQ classes for Java.

If you find that you cannot restart a failed instance, even though failover has occurred and the standby instance has become active, check to see whether applications connected locally to the failed instance have disconnected from the failed instance.

Locally connected applications must end or disconnect from a failed queue manager instance in order for the failed instance to be restarted. Any locally connected applications using shared bindings (which is the default setting) which hold on to a connection to a failed instance act to prevent the instance from being restarted.

If it is not possible to end the locally connected applications, or ensure that they disconnect when the local queue manager instance fails, consider using isolated bindings. Locally connected applications using isolated bindings do not prevent the local queue manager instance from being restarted, even if they do not disconnect.

## Channel and client reconnection:

Channel and client reconnection is an essential part of restoring message processing after a standby queue manager instance has become active.

Multi-instance queue manager instances are installed on servers with different network addresses. You need to configure IBM MQ channels and clients with connection information for all queue manager instances. When a standby takes over, clients and channels are automatically reconnected to the newly active queue manager instance at the new network address. Automatic client reconnect is not supported by IBM MQ classes for Java.

The design is different from the way high availability environments such as HA-CMP work. HA-CMP provides a virtual IP address for the cluster and transfer the address to the active server. IBM MQ reconnection does not change or reroute IP addresses. It works by reconnecting using the network addresses you have defined in channel definitions and client connections. As an administrator, you need to define the network addresses in channel definitions and client connections to all instances of any multi-instance queue manager. The best way to configure network addresses to a multi-instance queue manager depends on the connection:

### Queue manager channels

The CONNAME attribute of channels is a comma-separated list of connection names; for example, CONNAME('127.0.0.1(1234), 192.0.2.0(4321)'). The connections are tried in the order specified in the connection list until a connection is successfully established. If no connection is successful, the channel attempts to reconnect.

### Cluster channels

Typically, no additional configuration is required to make multi-instance queue managers work in a cluster.

If a queue manager connects to a repository queue manager, the repository discovers the network address of the queue manager. It refers to the CONNAME of the CLUSRCVR channel at the queue manager. On TCPIP, the queue manager automatically sets the CONNAME if you omit it, or configure it to blanks. When a standby instance takes over, its IP address replaces the IP address of the previous active instance as the CONNAME.

If it is necessary, you can manually configure CONNAME with the list of network addresses of the queue manager instances.

### Client connections

Client connections can use connection lists, or queue manager groups to select alternative connections. Clients need to be compiled to run with IBM WebSphere MQ Version 7.0.1 client libraries or better. They must be connected to at least a Version 7.0.1 queue manager.

When failover occurs, reconnection takes some time. The standby queue manager has to complete its startup. The clients that were connected to the failed queue manager have to detect the connection failure, and start a new client connection. If a new client connection selects the standby queue manager that has become newly active, then the client is reconnected to the same queue manager.

If the client is in the middle of an MQI call during the reconnection, it must tolerate an extended wait before the call completes.

If the failure takes place during a batch transfer on a message channel, the batch is rolled back and restarted.

Switching over is faster than failing over, and takes only as long as stopping one instance of the queue manager and starting another. For a queue manager with only few log records to replay, at best switchover might take of the order of a few seconds. To estimate how long failover takes, you need to add the time that it takes for the failure to be detected. At best the detection takes of the order of 10

seconds, and might be several minutes, depending on the network and the file system.

### **Application recovery:**

Application recovery is the automated continuation of application processing after failover. Application recovery following failover requires careful design. Some applications need to be aware failover has taken place.

The objective of application recovery is for the application to continue processing with only a short delay. Before continuing with new processing, the application must back out and resubmit the unit of work that it was processing during the failure.

A problem for application recovery is losing the context that is shared between the IBM MQ MQI client and the queue manager, and stored in the queue manager. The IBM MQ MQI client restores most of the context, but there are some parts of the context that cannot be reliably restored. The following sections describe some properties of application recovery and how they affect the recovery of applications connected to a multi-instance queue manager.

### **Transactional messaging**

From the perspective of delivering messages, failover does not change the persistent properties of IBM MQ messaging. If messages are persistent, and correctly managed within units of work, then messages are not lost during a failover.

From the perspective of transaction processing, transactions are either backed out or committed after failover.

Uncommitted transactions are rolled back. After failover, a re-connectable application receives a MQRC\_BACKED\_OUT reason code to indicate that the transaction has failed. It then needs to restart the transaction again.

Committed transactions are transactions that have reached the second phase of a two-phase commit, or single phase (message only) transactions that have begun MQCMIT.

If the queue manager is the transaction coordinator and MQCMIT has begun the second phase of its two-phase commit before the failure, the transaction successfully completes. The completion is under the control of the queue manager and continues when the queue manager is running again. In a reconnectable application, the MQCMIT call completes normally.

In a single phase commit, which involves only messages, a transaction that has started commit processing completes normally under the control of the queue manager once it is running again. In a reconnectable application, the MQCMIT completes normally.

Reconnectable clients can use single phase transactions under the control of the queue manager as the transaction coordinator. The extended transactional client does not support reconnection. If reconnection is requested when the transactional client connects, the connection succeeds, but without the ability to be reconnected. The connection behaves as if it is not reconnectable.

### **Application restart or resume**

Failover interrupts an application. After a failure an application can restart from the beginning, or it can resume processing following the interruption. The latter is called *automatic client reconnection*. Automatic client reconnect is not supported by IBM MQ classes for Java.

With an IBM MQ MQI client application, you can set a connection option to reconnect the client automatically. The options are MQCNO\_RECONNECT or MQCNO\_RECONNECT\_Q\_MGR. If no option is set, the client

does not try to reconnect automatically and the queue manager failure returns `MQRC_CONNECTION_BROKEN` to the client. You might design the client to try and start a new connection by issuing a new `MQCONN` or `MQCONNX` call.

Server programs have to be restarted; they cannot be automatically reconnected by the queue manager at the point they were processing when the queue manager or server failed. IBM MQ server programs are typically not restarted on the standby queue manager instance when a multi-instance queue manager instance fails.

You can automate an IBM MQ server program to restart on the standby server in two ways:

1. Package your server application as a queue manager service. It is restarted when the standby queue manager restarts.
2. Write your own failover logic, triggered for example, by the failover log message written by a standby queue manager instance when it starts. The application instance then needs to call `MQCONN` or `MQCONNX` after it starts, to create a connection to the queue manager.

### Detecting failover

Some applications do need to be aware of failover, others do not. Consider these two examples.

1. A messaging application that gets or receives messages over a messaging channel does not normally require the queue manager at the other end of the channel to be running; it is unlikely to be affected if the queue manager at the other end of the channel restarts on a standby instance.
2. An IBM MQ MQI client application processes persistent message input from one queue and puts persistent message responses onto another queue as part of a single unit of work: if it handles an `MQRC_BACKED_OUT` reason code from `MQPUT`, `MQGET`, or `MQCMIT` within sync point by restarting the unit of work, then no messages are lost. Additionally the application does not need to do any special processing to deal with a connection failure.

Suppose however, in the second example, that the application is browsing the queue to select the message to process by using the `MQGET` option, `MQGMO_MSG_UNDER_CURSOR`. Reconnection resets the browse cursor, and the `MQGET` call does not return the correct message. In this example, the application has to be aware failover has occurred. Additionally, before issuing another `MQGET` for the message under the cursor, the application must restore the browse cursor.

Losing the browse cursor is one example of how the application context changes following reconnection. Other cases are documented in “Recovery of an automatically reconnected client” on page 1125.

You have three alternative design patterns for IBM MQ MQI client applications following failover. Only one of them does not need to detect the failover.

### No reconnection

In this pattern, the application stops all processing on the current connection when the connection is broken. For the application to continue processing, it must establish a new connection with the queue manager. The application is entirely responsible for transferring any state information it requires to continue processing on the new connection. Existing client applications that reconnect with a queue manager after losing their connection are written in this way.

The client receives a reason code, such as `MQRC_CONNECTION_BROKEN`, or `MQRC_Q_MGR_NOT_AVAILABLE` from the next MQI call after the connection is lost. The application must discard all its IBM MQ state information, such as queue handles, and issue a new `MQCONN` or `MQCONNX` call to establish a new connection, and then reopen the IBM MQ objects it needs to process.

The default MQI behavior is for the queue manager connection handle to become unusable after a connection with the queue manager is lost. The default is equivalent to setting the `MQCNO_RECONNECT_DISABLED` option on `MQCONNX` to prevent application reconnection after failover.

## Failover tolerant

Write the application so it is unaffected by failover. Sometimes careful error handling is sufficient to deal with failover.

## Reconnection aware

Register an MQCBT\_EVENT\_HANDLER event handler with the queue manager. The event handler is posted with MQRC\_RECONNECTING when the client starts to try to reconnect to the server, and MQRC\_RECONNECTED after a successful reconnection. You can then run a routine to reestablish a predictable state so that the client application is able to continue processing.

## Recovery of an automatically reconnected client

Failover is an unexpected event, and for an automatically reconnected client to work as designed the consequences of reconnection must be predictable.

A major element of turning an unexpected failure into a predictable and reliable recovery is the use of transactions.

In the previous section, an example, 2 on page 1124, was given of an IBM MQ MQI client using a local transaction to coordinate MQGET and MQPUT. The client issues an MQCMIT or MQBACK call in response to a MQRC\_BACKED\_OUT error and then resubmits the backed out transaction. The queue manager failure causes the transaction to be backed out, and the behavior of the client application ensures no transactions, and no messages, are lost.

Not all program state is managed as part of a transaction, and therefore the consequences of reconnection become harder to understand. You need to know how reconnection changes the state of an IBM MQ MQI client in order to design your client application to survive queue manager failover.

You might decide to design your application without any special failover code, handling reconnection errors with the same logic as other errors. Alternatively, you might choose to recognize that reconnection requires special error processing, and register an event handler with IBM MQ to run a routine to handle failover. The routine might handle the reconnection processing itself, or set a flag to indicate to the main program thread that when it resumes processing it needs to perform recovery processing.

The IBM MQ MQI client environment is aware of failover itself, and restores as much context as it can, following reconnection, by storing some state information in the client, and issuing additional MQI calls on behalf of the client application to restore its IBM MQ state. For example, handles to objects that were open at the point of failure are restored, and temporary dynamic queues are opened with the same name. But there are changes that are unavoidable and you need your design to deal with these changes. The changes can be categorized into five kinds:

1. New, or previously undiagnosed errors, are returned from MQI calls until a consistent new context state is restored by the application program.

An example of receiving a new error is the return code MQRC\_CONTEXT\_NOT\_AVAILABLE when trying to pass context after saving context before the reconnection. The context cannot be restored after reconnection because the security context is not passed to an unauthorized client program. To do so would let a malicious application program obtain the security context.

Typically, applications handle common and predictable errors in a carefully designed way, and relegate uncommon errors to a generic error handler. The error handler might disconnect from IBM MQ and reconnect again, or even stop the program altogether. To improve continuity, you might need to deal with some errors in a different way.

2. Non-persistent messages might be lost.
3. Transactions are rolled back.
4. MQGET or MQPUT calls used outside a sync point might be interrupted with the possible loss of a message.

5. Timing induced errors, due to a prolonged wait in an MQI call.

Some details about lost context are listed in the following section.

- Non-persistent messages are discarded, unless put to a queue with the NPMCLASS(HIGH) option, and the queue manager failure did not interrupt the option of storing non-persistent messages on shutdown.
- A non-durable subscription is lost when a connection is broken. On reconnection, it is re-established. Consider using a durable subscription.
- The get-wait interval is recomputed; if its limit is exceeded it returns MQRC\_NO\_MSG\_AVAILABLE. Similarly, subscription expiry is recomputed to give the same overall expiry time.
- The position of the browse cursor in a queue is lost; it is typically reestablished before the first message.
  - MQGET calls that specify MQGMO\_BROWSE\_MSG\_UNDER\_CURSOR or MQGMO\_MSG\_UNDER\_CURSOR, fail with reason code MQRC\_NO\_MSG\_AVAILABLE.
  - Messages locked for browsing are unlocked.
  - Browse marked messages with handle scope are unmarked and can be browsed again.
  - Cooperatively browse marked messages are unmarked in most cases.
- Security context is lost. Attempts to use saved message context, such as putting a message with MQPMO\_PASS\_ALL\_CONTEXT fail with MQRC\_CONTEXT\_NOT\_AVAILABLE.
- Message tokens are lost. MQGET using a message token returns the reason code MQRC\_NO\_MSG\_AVAILABLE.

**Note:** *MsgId* and *CorrelId*, as they are part of the message, are preserved with the message during failover, and so MQGET using *MsgId* or *CorrelId* work as expected.

- Messages put on a queue under sync point in an uncommitted transaction are no longer available.
- Processing messages in a logical order, or in a message group, results in a return code of MQRC\_RECONNECT\_INCOMPATIBLE after reconnection.
- An MQI call might return MQRC\_RECONNECT\_FAILED rather than the more general MQRC\_CONNECTION\_BROKEN that clients typically receive today.
- Reconnection during an MQPUT call outside sync point returns MQRC\_CALL\_INTERRUPTED if the IBM MQ MQI client does not know if the message was delivered to the queue manager successfully. Reconnection during MQCMIT behaves similarly.
- MQRC\_CALL\_INTERRUPTED is returned - after a successful reconnect - if the IBM MQ MQI client has received no response from the queue manager to indicate the success or failure of
  - the delivery of a persistent message using an MQPUT call outside sync point.
  - the delivery of a persistent message or a message with default persistence using an MQPUT1 call outside sync point.
  - the commit of a transaction using an MQCMIT call. The response is only ever returned after a successful reconnect.
- Channels are restarted as new instances (they might also be different channels), and so no channel exit state is retained.
- Temporary dynamic queues are restored as part of the process of recovering reconnectable clients that had temporary dynamic queues open. No messages on a temporary dynamic queue are restored, but applications that had the queue open, or had remembered the name of the queue, are able to continue processing.

There is the possibility that if the queue is being used by an application other than the one that created it, that it might not be restored quickly enough to be present when it is next referenced. For example, if a client creates a temporary dynamic queue as a reply-to queue, and a reply message is to be placed on the queue by a channel, the queue might not be recovered in time. In this case, the channel would typically place the reply-to message on the dead letter queue.

If a reconnectable client application opens a temporary dynamic queue by name (because another application has already created it), then when reconnection occurs, the IBM MQ MQI client is unable to recreate the temporary dynamic queue because it does not have the model to create it from. In the MQI, only one application can open the temporary dynamic queue by model. Other applications that wish to use the temporary dynamic queue must use MQPUT1, or server bindings, or be able to try the reconnection again if it fails.

Only non-persistent messages might be put to a temporary dynamic queue, and these messages are lost during failover; this loss is true for messages being put to a temporary dynamic queue using MQPUT1 during reconnection. If failover occurs during the MQPUT1, the message might not be put, although the MQPUT1 succeeds. One workaround to this problem is to use permanent dynamic queues. Any server bindings application can open the temporary dynamic queue by name because it is not reconnectable.

### **Data recovery and high availability:**

High availability solutions using multi-instance queue managers must include a mechanism to recover data after a storage failure.

A multi-instance queue manager increases the availability of queue manager processes, but not the availability of other components, such as the file system, that the queue manager uses to store messages, and other information.

One way to make data highly available is to use networked resilient data storage. You can either build your own solution using a networked file system and resilient data storage, or you can buy an integrated solution. If you want to combine resilience with disaster recovery, then asynchronous disk replication, which permits disk replication over tens, or hundreds of kilometers, is available.

You can configure the way different IBM MQ directories are mapped to storage media, to make the best use of the media. For *multi-instance* queue managers there is an important distinction between two types of IBM MQ directories and files.

#### **Directories that must be shared between the instances of a queue manager.**

The information that must be shared between different instances of a queue manager is in two directories: the `qmgrs` and `logs` directories. The directories must be on a shared networked file system. You are advised to use a storage media that provides continuous high availability and excellent performance because the data is constantly changing as messages are created and deleted.

#### **Directories and files that do not *have* to be shared between instances of a queue manager.**

Some other directories do not have to be shared between different instances of a queue manager, and are quickly restored by means other than using a mirrored file system.

- IBM MQ executable files and the tools directory. Replace by reinstalling or by backing up and restoring from a backed up file archive.
- Configuration information that is modified for the installation as a whole. The configuration information is either managed by IBM MQ, such as the `mqs.ini` file on Windows, UNIX and Linux systems, or part of your own configuration management such as **MQSC** configuration scripts. Back up and restore using a file archive.
- Installation-wide output such as traces, error logs and FFDC files. The files are stored in the `errors` and `trace` subdirectories in the default data directory. The default data directory on UNIX and Linux systems is `/var/mqm`. On Windows the default data directory is the IBM MQ installation directory.

You can also use a backup queue manager to take regular media backups of a multi-instance queue manager using linear logging. A backup queue manager does not provide recovery that is as fast as from a mirrored file system, and it does not recover changes since the last backup. The backup queue manager mechanism is more appropriate for use in off-site disaster recovery scenarios than recovering a queue

manager after a localized storage failure.

### **Combining IBM MQ Availability solutions:**

Applications are using other IBM MQ capabilities to improve availability. Multi-instance queue managers complement other high availability capabilities.

### **IBM MQ Clusters increase queue availability**

You can increase queue availability by creating multiple definitions of a cluster queue; up to one of every queue on each manager in the cluster.

Suppose a member of the cluster fails and then a new message is sent to a cluster queue. Unless the message *has* to go to the queue manager that has failed, the message is sent to another running queue manager in the cluster that has a definition of the queue.

Although clusters greatly increase availability, there are two related failure scenarios that result in messages getting delayed. Building a cluster with multi-instance queue managers reduces the chance of a message being delayed.

### **Marooned messages**

If a queue manager in the cluster fails, no more messages that can be routed to other queue managers in the cluster are routed to the failed queue manager. Messages that have already been sent are marooned until the failed queue manager is restarted.

### **Affinities**

Affinity is the term used to describe information shared between two otherwise separate computations. For example, an affinity exists between an application sending a request message to a server and the same application expecting to process the reply. Another example would be a sequence of messages, the processing of each message depending on the previous messages.

If you send messages to clustered queues you need to consider affinities. Do you need to send successive messages to the same queue manager, or can each message go to any member of the cluster?

If you do need to send messages to the same queue manager in the cluster and it fails, the messages wait in the transmission queue of the sender until the failed cluster queue manager is running again.

If the cluster is configured with multi-instance queue managers the delay waiting for the failed queue manager to restart is limited to the order of a minute or so while the standby takes over. When the standby is running, marooned messages resume processing, channels to the newly activated queue manager instance are started, and the messages that were waiting in transmission queues start flowing.

A possible way to configure a cluster to overcome messages being delayed by a failed queue manager, is to deploy two different queue managers to each server in the cluster, and arrange for one to be the active and one to be the standby instance of the different queue managers. This is an active-standby configuration, and it increases the availability of the cluster.

As well as having the benefits of reduced administration and increased scalability, clusters continue to provide additional elements of availability to complement multi-instance queue managers. Clusters protect against other types of failure that affect both the active and standby instances of a queue manager.

### **Uninterrupted service**

A cluster provides an uninterrupted service. New messages received by the cluster are sent to active queue managers to be processed. Do not rely on a multi-instance queue manager to



provide an uninterrupted service because it takes time for the standby queue manager to detect the failure and complete its startup, for its channels to be reconnected, and for failed batches of messages to be resubmitted.

### **Localized outage**

There are practical limitations to how far apart the active, standby, and file system servers can be, as they need to interact at millisecond speeds to deliver acceptable performance.

Clustered queue managers require interaction speeds of the order of many seconds, and can be geographically dispersed anywhere in the world.

### **Operational error**

By using two different mechanisms to increase availability you reduce the chances that an operational error, such as a human error, compromises your availability efforts.

### **Queue sharing groups increase message processing availability**

Queue sharing groups, provided only on z/OS, allow a group of queue managers to share servicing a queue. If one queue manager fails, the other queue managers continue to process all the messages on the queue. Multi-instance queue managers are not supported on z/OS and complement queue sharing groups only as part of a wider messaging architecture.

### **IBM MQ Clients increase application availability**

IBM MQ MQI client programs can connect to different queue managers in a queue manager group based on queue manager availability, connection weightings, and affinities. By running an application on a different machine from the one on which the queue manager is running, you can improve the overall availability of a solution as long as there is a way to reconnect the application if the queue manager instance it is connected to fails.

Queue manager groups are used to increase client availability by uncoupling a client from a queue manager that is stopped, and load balancing client connections across a group of queue managers, rather like an IP sprayer. The client application must have no affinities with the failed queue manager, such as a dependency on a particular queue, or it cannot resume processing.

Automatic client reconnection and multi-instance queue managers increase client availability by resolving some affinity problems. Automatic client reconnect is not supported by IBM MQ classes for Java.

You can set the MQCNO option `MQCNO_RECONNECT_Q_MGR`, to force a client to reconnect to the same queue manager:

1. If the previously connected single instance queue manager is not running the connection is tried again until the queue manager is running again.
2. If the queue manager is configured as a multi-instance queue manager, then the client reconnects to whichever instance is active.

By automatically reconnecting to the same queue manager, much of the state information the queue manager was holding on behalf of the client, such as the queues it had open and the topic it was subscribed to, are restored. If the client had opened a dynamic reply-to queue to receive a reply to a request, the connection to the reply-to queue is restored too.

## Making sure that messages are not lost (logging)

IBM MQ records all significant changes to the persistent data controlled by the queue manager in a recovery log.

This includes creating and deleting objects, persistent message updates, transaction states, changes to object attributes, and channel activities. The log contains the information you need to recover all updates to message queues by:

- Keeping records of queue manager changes
- Keeping records of queue updates for use by the restart process
- Enabling you to restore data after a hardware or software failure

However, IBM MQ also relies on the disk system hosting its files, including log files. If the disk system is itself unreliable, information, including log information, can still be lost.

### What logs look like

Logs consist of primary and secondary files, and a control file. You define the number and size of log files and where they are stored in the file system.

An IBM MQ log consists of two components:

1. One or more files of log data.
2. A log control file

A file of log data is also known as a log extent.

There are a number of log extents that contain the data being recorded. You can define the number and size (as explained in “Log defaults for IBM MQ” on page 777), or take the system default of three primary and two secondary extents.

Each of the three primary and two secondary extents defaults to 16 MB.

When you create a queue manager, the number of log extents pre-allocated is the number of *primary* log extents allocated. If you do not specify a number, the default value is used.

IBM MQ uses two types of logging:

- Circular
- Linear

The number of log extents used with linear logging can be very large, depending on the frequency of your media image recording.

See “Types of logging” on page 1131 for more information.

In IBM MQ for Windows, if you have not changed the log path, log extents are created under the directory:

```
C:\ProgramData\IBM\MQ\log\<QMGrName>
```

In IBM MQ for UNIX and Linux systems, if you have not changed the log path, log extents are created under the directory:

```
/var/mqm/log/<QMGrName>
```

IBM MQ starts with these primary log extents, but if the primary log space is not sufficient, it allocates *secondary* log extents. It does this dynamically and removes them when the demand for log space reduces. By default, up to two secondary log extents can be allocated. You can change this default allocation, as described in “Changing IBM MQ and queue manager configuration information” on page 755.

## The log control file:

The log control file contains information needed to describe the state of log extents, such as their size and location and the name of the next available extent.

**Important:** The log control file is for internal queue manager use only.

The queue manager keeps control data associated with the state of the recovery log in the log control file and you must not modify the contents of the log control file.

The log control file is in the log path and is called `amqh1ctl.lfh`. When backing up or restoring your queue manager, ensure that the log control file is backed up and restored, along with your log extents.

## Types of logging

In IBM MQ there are two ways of maintaining records of queue manager activities: circular logging and linear logging.

### Circular logging

Use circular logging if all you want is restart recovery, using the log to roll back transactions that were in progress when the system stopped.

Circular logging keeps all restart data in a ring of log files. Logging fills the first file in the ring, then moves on to the next, and so on, until all the files are full. It then goes back to the first file in the ring and starts again. This continues as long as the product is in use, and has the advantage that you never run out of log files.

IBM MQ keeps the log entries required to restart the queue manager without loss of data until they are no longer required to ensure queue manager data recovery. The mechanism for releasing log files for reuse is described in “Restart recovery” on page 1133.

### Linear logging

Use linear logging if you want both restart recovery and media recovery (re-creating lost or damaged data by replaying the contents of the log). Linear logging keeps the log data in a continuous sequence of log files.

Log files can optionally be:

- Reused, but only when they are no longer needed for either restart recovery or media recovery.
- Manually archived for longer term storage and analysis.

The frequency of media images determines when linear log files can be reused, and is a major factor in how much disk space must be available for linear log files.

You can configure the queue manager to automatically take periodic media images, based either upon time or log usage, or you can schedule media images manually.

Your administrator decides what policy to implement, and the implications on disk space usage. Log files needed for restart recovery must always be available, whereas log files needed only for media recovery can be archived to longer term storage, for example, tape.

If your administrator enables automatic log management and automatic media images, linear logging behaves in a similar way to a very large circular log, but with the improved redundancy against media failure enabled by media recovery.

## Active log

There are a number of files that are said to be *active* in both linear and circular logging. The active log is the maximum amount of log space, whether you are using circular or linear logging, that might be referenced by restart recovery.

The number of active log files is usually less than the number of primary log files as defined in the configuration files. (See “Calculating the size of the log” on page 1136 for information about defining the number.)

Note, that the active log space does not include the space required for media recovery, and that the number of log files used with linear logging can be very large, depending on your message flow and the frequency of media images.

## Inactive log

When a log file is no longer needed for restart recovery it becomes *inactive*. Log files that are not required for either restart recovery, or media recovery, can be considered as superfluous log files.

When using automatic log management, the queue manager controls the processing of these superfluous log files. If you have selected manual log management, it becomes the responsibility of your administrator to manage (for example, delete and archive) superfluous log files if they are no longer of interest to your operation.

Refer to “Managing logs” on page 1140 for further information about the disposition of log files.

## Secondary log files

Although secondary log files are defined for linear logging, they are not used in normal operation. If a situation arises when, probably due to long-lived transactions, it is not possible to free a file from the active pool because it might still be required for a restart, secondary files are formatted and added to the active log file pool.

If the number of secondary files available is used up, requests for most further operations requiring log activity will be refused with an MQRC\_RESOURCE\_PROBLEM return code being returned to the application, and any long running transactions will be considered for asynchronous rollback.

**Attention:** Both types of logging can cope with unexpected loss of power, assuming that there is no hardware failure.

## Restart recovery

Both circular logging and linear logging queue managers support restart recovery. Regardless of how abruptly the previous instance of the queue manager terminates (for example a power outage) upon restart the queue manager restores its persistent state to the correct transactional state at the point of termination.

Restart recovery depends upon disk integrity being maintained. Similarly, the operating system should ensure disk integrity regardless of how abruptly an operating system termination might occur.

In the highly unusual event that disk integrity is not maintained then linear logging (and media recovery) provides some further redundancy and recoverability options. With increasingly common technology, such as RAID, it is increasingly rare to suffer disk integrity issues and many enterprises configure circular logging and use only restart recovery.

IBM MQ is designed as a classic Write Ahead Logging resource manager. Persistent updates to message queues happen in two stages:

1. Log records representing the update are written reliably to the recovery log
2. The queue file or buffers are updated in a manner that is the most efficient for your system, but not necessarily consistently.

The log files can thus become more up to date than the underlying queue buffer and file state.

If this situation was allowed to continue unabated, then a very large volume of log replay would be required to make the queue state consistent following a crash recovery.

IBM MQ uses checkpoints in order to limit the volume of log replay required following a crash recovery. The key event that controls whether a log file is termed active or not is a checkpoint.

An IBM MQ checkpoint is a point:

- Of consistency between the recovery log and object files.
- That identifies a place in the log, from which forward replay of subsequent log records is guaranteed to restore the queue to the correct logical state at the time the queue manager might have ended.

During a checkpoint, IBM MQ flushes older updates to the queues files, as required, in order to limit the volume of log records that need to be replayed to bring the queues back to a consistent state following a crash recovery.

The most recent complete checkpoint marks a point in the log from which replay must be performed during crash recovery. The frequency of checkpoint is thus a trade-off between the overhead of recording checkpoints, and the improvement in potential recovery time implied by those checkpoints.

The position in the log of the start of the most recent complete checkpoint is one of the key factors in determining whether a log file is active or inactive. The other key factor is the position in the log of the first log record relating to the first persistent update made by a current active transaction.

If a new checkpoint is recorded in the second, or later, log file and no current transaction refers to a log record in the first log file, the first log file become inactive. In the case of circular logging the first log file is now ready to be reused. In the case of linear logging the first log file will typically still be required for media recovery.

If you configure either circular logging or automatic log management the queue manager will manage the inactive log files. If you configure linear logging with manual log management it becomes an administrative task to manage the inactive files according to the requirements of your operation.

IBM MQ generates checkpoints automatically. They are taken when:

- The queue manager starts
- At shutdown
- When logging space is running low
- **distributed** After 50,000 operations have been logged since the previous checkpoint was taken
- **z/OS** For z/OS, the LOGLOAD setting controls how many operations are in a checkpoint.

When IBM MQ restarts, it finds the latest checkpoint record in the log. This information is held in the checkpoint file that is updated at the end of every checkpoint. All the operations that have taken place since the checkpoint are replayed forward. This is known as the replay phase.

The replay phase brings the queues back to the logical state they were in before the system failure or shutdown. During the replay phase a list is created of the transactions that were in-flight when the system failure or shutdown occurred.

**distributed** Messages AMQ7229 and AMQ7230 are issued to indicate the progression of the replay phase.

In order to know which operations to back out or commit, IBM MQ accesses each active log record associated with an in-flight transaction. This is known as the recovery phase.

**distributed** Messages AMQ7231, AMQ7232 and AMQ7234 are issued to indicate the progression of the recovery phase.

Once all the necessary log records have been accessed during the recovery phase, each active transaction is in turn resolved and each operation associated with the transaction will be either backed out or committed. This is known as the resolution phase.

**distributed** Message AMQ7233 is issued to indicate the progression of the resolution phase.

**z/OS** On z/OS, restart processing is made up of various phases.


1. The recovery log range is established, based on the media recovery required for the page sets and the oldest log record that is required for backing out units of work and obtaining locks for in-doubt units of work.
2. Once the log range has been determined, forward log reading is carried out to bring the page sets up to the latest state, and also to lock any messages that are related to in-doubt or in-flight units of work.
3. When forward log reading has been completed the logs are read backwards to backout any units of work that were in-flight or in-backout at the time of failure.

**z/OS** An example of the messages you might see:

```
CSQR001I +MQOX RESTART INITIATED
CSQR003I +MQOX RESTART - PRIOR CHECKPOINT RBA=00000001E48C0A5E
CSQR004I +MQOX RESTART - UR COUNTS - 806
IN COMMIT=0, INDOUBT=0, INFLIGHT=0, IN BACKOUT=0
CSQR030I +MQOX Forward recovery log range 815
from RBA=00000001E45FF7AD to RBA=00000001E48C1882
CSQR005I +MQOX RESTART - FORWARD RECOVERY COMPLETE - 816
IN COMMIT=0, INDOUBT=0
CSQR032I +MQOX Backward recovery log range 817
from RBA=00000001E48C1882 to RBA=00000001E48C1882
CSQR006I +MQOX RESTART - BACKWARD RECOVERY COMPLETE - 818
INFLIGHT=0, IN BACKOUT=0
CSQR002I +MQOX RESTART COMPLETED
```

**Note:** If there is a large amount of log to be read, messages CSQR031I (forward recovery) and CSQR033I (backwards recovery) are issued periodically to show the progression.

In Figure 160, all records before the latest checkpoint, Checkpoint 2, are no longer needed by IBM MQ. The queues can be recovered from the checkpoint information and any later log entries. For circular logging, any freed files before the checkpoint can be reused. For a linear log, the freed log files no longer need to be accessed for normal operation and become inactive. In the example, the queue head pointer is moved to point at the latest checkpoint, Checkpoint 2, which then becomes the new queue head, Head 2. Log File 1 can now be reused.



*Figure 160. Checkpointing.* For simplicity, only the ends of the log files are shown.

### **Checkpointing with long-running transactions:**

How a long-running transaction affects reuse of log files.

Figure 161 on page 1136 shows how a long-running transaction affects reuse of log files. In the example, a long-running transaction has made an entry to the log, shown as LR 1, after the first checkpoint shown. The transaction does not complete (at point LR 2) until after the third checkpoint. All the log information from LR 1 onwards is retained to allow recovery of that transaction, if necessary, until it has completed.

After the long-running transaction has completed, at LR 2, the head of the log logically moves to Checkpoint 3, the latest logged checkpoint. The files containing log records before Checkpoint 3, Head 2, are no longer needed. If you are using circular logging, the space can be reused.

If the primary log files are completely full before the long-running transaction completes, secondary log files might be used to avoid the logs getting full.

Activities which are entirely under the control of the queue manager, for example checkpointing, are scheduled to try and keep the activity within the primary log.

However, when secondary log space is required to support behavior outside of the control of the queue manager (for example the duration of one of your transactions) the queue manager tries using any defined secondary log space, to allow that activity to complete.

If that activity does not complete by the time 80% of the total log space is in use, the queue manager initiates action to reclaim log space, regardless of the fact that this has an impact on the application.

When the log head is moved and you are using circular logging, the primary log files might become eligible for reuse and the logger, after filling the current file, reuses the first primary file available to it. If you are using linear logging, the log head is still moved down the active pool and the first file becomes inactive. A new primary file is formatted and added to the bottom of the pool in readiness for future logging activities.

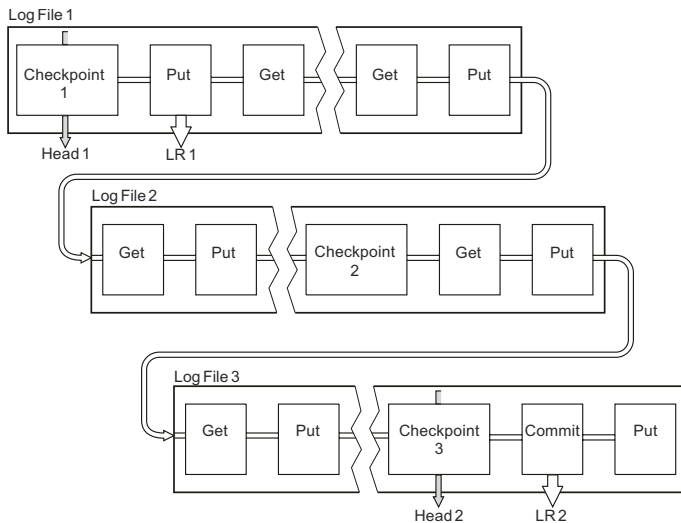


Figure 161. Checkpointing with a long-running transaction. For simplicity, only the ends of the log files are shown.

## Calculating the size of the log

Estimating the size of log a queue manager needs.

After deciding whether the queue manager uses circular or linear logging, you need to estimate the size of the Active log that the queue manager needs. The size of the active log is determined by the following log configuration parameters:

### LogFilePages

The size of each primary and secondary log file in units of 4K pages

### LogPrimaryFiles

The number of preallocated primary log files

### LogSecondaryFiles

The number of secondary log files that can be created for use when the primary log files are becoming full

### Notes:

1. You can change the number of primary and secondary log files each time the queue manager starts, although you might not notice the effect of the change you make to the secondary logs immediately.
2. You cannot change the log file size; you must determine it **before** creating the queue manager.



3. The number of primary log files and the log file size determine the amount of log space that is preallocated when the queue manager is created.
4. The total number of primary and secondary log files cannot exceed 511 on UNIX and Linux systems, or 255 on Windows, which in the presence of long-running transactions, limits the maximum amount of log space available to the queue manager for restart recovery. The amount of log space the queue manager might need for media recovery does not share this limit.
5. When *circular* logging is being used, the queue manager reuses primary and secondary log space. The queue manager will, up to a limit, allocate a secondary log file when a log file becomes full, and the next primary log file in the sequence is not available.  
See “How large should I make my active log?” for information on the number of logs you need to allocate. The primary log extents are used in sequence and that sequence does not change.  
For example, if you have three primary logs 0, 1, and 2, the order of use is 0,1,2 followed by 1,2,0, 2,0,1, back to 0,1,2 and so on. Any secondary logs you have allocated are interspersed as required.
6. Primary log files are made available for reuse during a checkpoint. The queue manager takes both the primary and secondary log space in to consideration before taking a checkpoint because the amount of log space is running low.

See “Log defaults for IBM MQ” on page 777 for more information.

### How large should I make my active log?:

Estimating the size of active log a queue manager needs.

The size of the active log is limited by:

$$\text{logsize} = (\text{primaryfiles} + \text{secondaryfiles}) * \text{logfilepages} * 4096$$

The log should be large enough to cope with your longest running transaction running when the queue manager is writing the maximum amount of data per second to disk.

If your longest running transaction runs for N seconds, and the maximum amount of data per second written to disk by the queue manager is B bytes per second in the log, your log should be at least:

$$\text{logsize} \geq 2 * (N+1) * B$$

The queue manager is likely to be writing the maximum amount of data per second to disk when you are running at peak workload, or it might be when you are recording media images.

If a transaction runs for so long that the log extent containing its first log record is not contained within the active log, the queue manager rolls back active transactions one at a time, starting with the transaction with the oldest log record.

The queue manager needs to make old log extents inactive before the maximum number of primary and secondary files are being used, and the queue manager needs to allocate another log extent.

Decide how long you want your longest running transaction to run, before the queue manager is allowed to roll it back. Your longest running transaction might be waiting for slow network traffic or, in the case of a poorly designed transaction, waiting for user input.

You can investigate how long your longest running transaction runs for, by issuing the following **runmqsc** command:

```
DISPLAY CONN(*) UOWLOGDA UOWLOGTI
```

Issuing the **dspmqrn -a** command, shows all the XA and non XA commands in all states.

Issuing this command lists the date and time that the first log record was written for all of your current transactions.

**Attention:** For the purposes of calculating the log size, it is the time since the first log record was written that matters, not the time since the application or transaction started. Round up the length of your longest running transaction to the nearest second. This is because of optimizations in the queue manager.

The first log record can be written long after the application started, if the application begins by, for example, issuing an MQGET call that waits for a length of time before actually getting a message.

By reviewing the maximum observed date and time output from the

```
DISPLAY CONN(*) UOWLOGDA UOWLOGTI
```

command you issued originally, from the current date and time, you can estimate how long your longest running transaction runs.

Ensure you run this **runmqsc** command repeatedly while your longest running transactions are running at peak workload so that you do not underestimate the length of your longest running transaction.

In IBM MQ Version 8.0 use the operating system tools, for example, **iostat** on UNIX platforms.

For example, if the operating system tools show that the logical bytes per second written to the log is approx 12 MB per second, using:

```
DISPLAY CONN(*) UOWLOGDA UOWLOGTI
```

shows that the longest running transaction was:

```
CONN(57E14F6820700069)
EXTCONN(414D51436D61726B2020202020202020)
TYPE(CONN)
APPLTAG(msginteg_r) UOWLOGDA(2016-09-20)
UOWLOGTI(16.44.14)
```

As the current date and time was 2016-09-20 16.44.19, this transaction had been running for 5 seconds. However, you require tolerating transactions running for 10 seconds before the queue manager rolls them back. So your log size should be:

$$2 * (10 + 1) * 12 = 264 \text{ MB}$$

.

Using the default **LogFilePages** which is 4096, you need to make sure that the sum of your **LogPrimaryFiles** and **LogSecondaryFiles** is at least 17:

$$264 \text{ MB} < 17 * 4096 * 4096$$

If you size your log so that your expected workload runs within the primary files:

- The secondary files provide some contingency in case additional log space is needed.
- Circular logging always using preallocated primary files, which is marginally faster than allocating and deallocating secondary files.
- The queue manager uses only the space remaining in the primary files to calculate when to take the next checkpoint.

Therefore, in the preceding example, set the following values:

- **LogFilePages** = 4096
- **LogPrimaryFiles** = 17
- **LogSecondaryFiles** = 5

Note the following:

- In this example, 5 secondaries is more than 20 per cent of the active log space.

You should be aware that the queue manager takes action to reduce log space usage when more than 80 per cent of the total log space is in use.

- Perform the same calculation whether you are using linear or circular logging.

It makes no difference whether you are calculating the size of a linear or circular active log, as the concept of the active log means the same in both linear logging and circular logging.

- The log extents needed for media recovery only are not within the active log, and are therefore not counted in the number of primary and secondary files.

### **What happens if I make my log too small?:**

Points you need to consider when estimating the minimum size of the log.

If you make your log too small:

- Long running transactions will get backed out.
- You might get FFDCs informing you that the log is too small.
- The next checkpoint wants to start before the previous one has ended.

**Important:** No matter how inaccurately you estimate the size of your log, data integrity is maintained.

See “Restart recovery” on page 1133 for an explanation of checkpoints. If the amount of log space left in the active log extents is becoming short, the queue manager schedules checkpoints more frequently.

A checkpoint takes some amount of time; it is not instantaneous. The more data that needs to be recorded in the checkpoint, the longer the checkpoint takes. If the log is small checkpoints can overlap, meaning that the next checkpoint is requested before the previous checkpoint has ended. If this happens error messages are written.

If long running transactions get backed out, or checkpoints overlap, the queue manager continues processing the workload. Short-lived transactions continue running as normal.

However, the queue manager is not running optimally and performance might be degraded. You should restart the queue manager with sufficient log space.

### **What happens if I make my log too large?:**

Points you need to consider when estimating the maximum size of the log.

If you make your log too large:

- You might increase the time taken for an emergency restart, although this is unlikely.
- You are using unnecessary disk space.
- Very long running transaction are tolerated.

**Important:** No matter how inaccurately you estimate the size of your log, data integrity is maintained.

See “Restart recovery” on page 1133 for a description of how the queue manager reads the log on restart. The queue manager replays the log from the last checkpoint, and then resolves all transactions that were active when the queue manager ended.

To resolve a transaction, the queue manager reads back all the log records associated with that transaction. These log records might predate the last checkpoint.

By allocating the queue manager a very large log, you are giving the queue manager permission to read every log record in the log on restart, although usually the queue manager does not have to do this. Potentially, in the unlikely event that this happens, that process could take a long time.

If checkpointing had unexpectedly stopped before the queue manager ended, that dramatically increases the restart time for a queue manager with a large log. Limiting the size of the log limits the emergency restart time.

To avoid these problems you should ensure that:

- Your workload can comfortably fit into a log that is not excessively large.
- You avoid long running transactions.

## Managing logs

You need to manage linear logs yourself. However, circular logs are nearly self-managing, but sometimes need intervention to resolve space problems.

On circular logging, the queue manager reclaims freed space in the log files. This activity is not apparent to the user, and you do not usually see the amount of disk space used reduce, because the space allocated is quickly reused.

On linear logging, the log might fill if a checkpoint has not been taken for a long time, or if a long-running transaction wrote a log record a long time ago. The queue manager tries to take checkpoints often enough to avoid the first problem.

**distributed** If the log fills, message AMQ7463 is issued. In addition, if the log fills because a long-running transaction has prevented the space being released, message AMQ7465 is issued.

Of the log records, only those written since the start of the last complete checkpoint, and those written by any active transactions, are needed to restart the queue manager.

Over time, the oldest log records written become unnecessary for restarting the queue manager.

When a long-running transaction is detected, activity is scheduled to asynchronously rollback that transaction. If, for some unexpected reason, that asynchronous rollback were to fail, some MQI calls return MQRC\_RESOURCE\_PROBLEM in that situation.

Note that space is reserved to commit or roll back all in-flight transactions, so **MQCMIT** or **MQBACK** should not fail.

The queue manager rolls back transactions that run for a long duration. An application that has a transaction is rolled back in this way cannot perform subsequent **MQPUT** or **MQGET** operations specifying sync point under the same transaction.

However, transactions ended manually start a new log. Note, that whereas new log space is allocated immediately, log space that has been released takes a finite time to be freed up.

An attempt to put or get a message under sync point in this state returns MQRC\_BACKED\_OUT. The application can then issue **MQCMIT**, which returns MQRC\_BACKED\_OUT, or **MQBACK** and start a new transaction. When the transaction consuming too much log space has been rolled back, the log space is released and the queue manager continues to operate normally.

### **What happens when a disk gets full:**

The queue manager logging component can cope with a full disk, and with full log files. If the disk containing the log fills, the queue manager issues message AMQ6709 and an error record is taken.

The log files are created at their fixed size, rather than being extended as log records are written to them. This means that IBM MQ can run out of disk space only when it is creating a new file; it cannot run out of space when it is writing a record to the log. IBM MQ always knows how much space is available in the existing log files, and manages the space within the files accordingly.

If you fill the drive containing the log files, you might be able to free some disk space. If you are using a linear log, there might be some inactive log files in the log directory, and you can copy these files to another drive or device.

Circular logging returns a resource problem.

If you still run out of space, check that the configuration of the log in the queue manager configuration file is correct. You might be able to reduce the number of primary or secondary log files so that the log does not outgrow the available space.

You cannot alter the size of the log files for an existing queue manager. The queue manager requires that all log extents are the same size.

### **Managing log files:**

Allocate sufficient space for your log files. For linear logging, you can delete old log files when they are no longer required.

### **Information specific to circular logging**

If you are using circular logging, ensure that there is sufficient space to hold the log files when you configure your system (see “Log defaults for IBM MQ” on page 777 and “Queue manager logs” on page 786 ). The amount of disk space used by the log does not increase beyond the configured size, including space for secondary files to be created when required.

### **Information specific to linear logging**

If you are using a linear log, the log files are added continually as data is logged, and the amount of disk space used increases with time. If the rate of data being logged is high, disk space is used rapidly by new log files.

Over time, the older log files for a linear log are no longer required to restart the queue manager or to perform media recovery of any damaged objects. The following are methods for determining which log files are still required:

#### **Logger event messages**

When a significant event occurs, for example a record media image, logger event messages are generated. The contents of logger event messages specify the log files that are still required for queue manager restart, and media recovery. For more information about logger event messages, see *Logger events*

#### **Queue manager status**

Running the MQSC command, DISPLAY QMSTATUS, or the PCF command, Inquire Queue Manager Status, returns queue manager information, including details of the required log files. For more information about MQSC commands, see *Script (MQSC) Commands*, and for information about PCF commands, see *Automating administration tasks*.

### Queue manager messages

Periodically, the queue manager issues a pair of messages to indicate which of the log files are needed:

- Message AMQ7467 gives the name of the oldest log file required to restart the queue manager. This log file and all newer log files must be available during queue manager restart.
- Message AMQ7468 gives the name of the oldest log file needed for media recovery.

To determine "older" and "newer" log files, use the log file number rather than the modification times applied by the file system.

### Information applicable to both types of logging

Only log files required for queue manager restart, active log files, are required to be online. Inactive log files can be copied to an archive medium such as tape for disaster recovery, and removed from the log directory. Inactive log files that are not required for media recovery can be considered as superfluous log files. You can delete superfluous log files if they are no longer of interest to your operation.

If any log file that is needed cannot be found, operator message AMQ6767 is issued. Make the log file, and all subsequent log files, available to the queue manager and try the operation again.

**Note:** When performing media recovery, all the required log files must be available in the log file directory at the same time. Make sure that you take regular media images of any objects you might want to recover to avoid running out of disk space to hold all the required log files.

For example, to take a media image of all your objects in your queue manager, run the **rcdmqimg** command as shown in the following examples:

#### On Windows

```
rcdmqimg -m QMNAME -t all *
```

#### On UNIX and Linux

```
rcdmqimg -m QMNAME -t all "*"
```

Running **rcdmqimg** moves the media log sequence number (LSN) forwards. For further details on log sequence numbers, see "Dumping the contents of the log using the **dmpmqlog** command" on page 1146. **rcdmqimg** does not run automatically, therefore must be run manually or from an automatic task you have created. For more information about this command, see **rcdmqimg** and **dmpmqlog**.

**Note:** Messages AMQ7467 and AMQ7468 can also be issued at the time of running the **rcdmqimg** command.

*Determining superfluous log files - linear logging only:*

For circular logging, never delete data from the log directory. When managing linear log files, it is important to be sure which files can be deleted or archived. This information will assist you in making this decision.

Do not use the file system's modification times to determine "older" log files. Use only the log file number. The queue manager's use of log files follows complex rules, including pre-allocation and formatting of log files before they are needed. You might see log files with modification times that would be misleading if you try to use these times to determine relative age.

To determine the oldest log file needed, there are three places available to you to use:

- The DISPLAY QMSTATUS command
- Logger event messages and, finally
- Error log messages

For the DISPLAY QMSTATUS command, to determine the oldest log extent needed to:

- Restart the queue manager, issue the command DISPLAY QMSTATUS RECLLOG.
- Perform media recovery, issue the command DISPLAY QMSTATUS MEDIALOG.

In general a lower log file number implies an older log. Unless you have a very high log file turnover, of the order of 3000 log files per day for 10 years, you do not need to cater for the number wrapping at 9 999 999. In this case, you can archive any log file with a number less than the RECLLOG value, and you can delete any log file with a number less than both the RECLLOG and MEDIALOG values.

**Attention:** The log file wraps, so the next number after 9 999 999 is zero.

*Log file location:*

When choosing a location for your log files, remember that operation is severely affected if IBM MQ fails to format a new log because of lack of disk space.

If you are using a circular log, ensure that there is sufficient space on the drive for at least the configured primary log files. Also leave space for at least one secondary log file, which is needed if the log has to grow.

If you are using a linear log, allow considerably more space; the space consumed by the log increases continuously as data is logged.

You should place the log files on a separate disk drive from the queue manager data.

Data integrity on this device is paramount - you should allow for built in redundancy.

It might also be possible to place the log files on multiple disk drives in a mirrored arrangement. This protects against failure of the drive containing the log. Without mirroring, you could be forced to go back to the last backup of your IBM MQ system.

## Using the log for recovery

Using logs to recover from failures.

There are several ways that your data can be damaged. IBM MQ helps you to recover from:

- A damaged data object
- A power loss in the system
- A communications failure

This section looks at how the logs are used to recover from these problems.

### Recovering from power loss or communications failures:

IBM MQ can recover from both communications failures and loss of power. In addition, it can sometimes recover from other types of problem, such as inadvertent deletion of a file.

In the case of a communications failure, persistent messages remain on queues until they are removed by a receiving application. If the message is being transmitted, it remains on the transmission queue until it can be successfully transmitted. To recover from a communications failure, you can usually restart the channels using the link that failed.

If you lose power, when the queue manager is restarted IBM MQ restores the queues to their committed state at the time of the failure. This ensures that no persistent messages are lost. Nonpersistent messages are discarded; they do not survive when IBM MQ stops abruptly.

### Recovering damaged objects:

There are ways in which an IBM MQ object can become unusable, for example because of inadvertent damage. You must then recover either your complete system or some part of it. The action required depends on when the damage is detected, whether the log method selected supports media recovery, and which objects are damaged.

### Media recovery

Media recovery re-creates objects from information recorded in a linear log. For example, if an object file is inadvertently deleted, or becomes unusable for some other reason, media recovery can re-create it. The information in the log required for media recovery of an object is called a *media image*.

A media image is a sequence of log records containing an image of an object from which the object itself can be re-created.

The first log record required to re-create an object is known as its *media recovery record* ; it is the start of the latest media image for the object. The media recovery record of each object is one of the pieces of information recorded during a checkpoint.

When an object is re-created from its media image, it is also necessary to replay any log records describing updates performed on the object since the last image was taken.

Consider, for example, a local queue that has an image of the queue object taken before a persistent message is put onto the queue. In order to re-create the latest image of the object, it is necessary to replay the log entries recording the putting of the message to the queue, in addition to replaying the image itself.

When an object is created, the log records written contain enough information to completely re-create the object. These records make up the first media image of the object. Then, at each shutdown, the queue manager records media images automatically as follows:



- Images of all process objects and queues that are not local
- Images of empty local queues

Media images can also be recorded manually using the **rctdmqimg** command, described in `rctdmqimg`. This command writes a media image of the IBM MQ object.

When a media image has been written, only the logs that hold the media image, and all the logs created after this time, are required to re-create damaged objects. The benefit of creating media images depends on such factors as the amount of free storage available, and the speed at which log files are created.

### Recovering from media images

IBM MQ automatically recovers some objects from their media image if it finds that they are corrupted or damaged. In particular, recovery applies to objects found to be damaged during the normal queue manager startup. If any transaction was incomplete when the queue manager last shut down, any queue affected is also recovered automatically in order to complete the startup operation.

You must recover other objects manually, using the **rcrmqobj** command, which replays the records in the log to re-create the IBM MQ object. The object is re-created from its latest image found in the log, together with all applicable log events between the time the image was saved and the time the re-create command was issued. If an IBM MQ object becomes damaged, the only valid actions that can be performed are either to delete it or to re-create it by this method. Nonpersistent messages cannot be recovered in this way.

See `rcrmqobj` for further details of the **rcrmqobj** command.

The log file containing the media recovery record, and all subsequent log files, must be available in the log file directory when attempting media recovery of an object. If a required file cannot be found, operator message AMQ6767 is issued and the media recovery operation fails. If you do not take regular media images of the objects that you want to re-create, you might have insufficient disk space to hold all the log files required to re-create an object.

### Recovering damaged objects during startup

If the queue manager discovers a damaged object during startup, the action it takes depends on the type of object and whether the queue manager is configured to support media recovery.

If the queue manager object is damaged, the queue manager cannot start unless it can recover the object. If the queue manager is configured with a linear log, and thus supports media recovery, IBM MQ automatically tries to re-create the queue manager object from its media images. If the log method selected does not support media recovery, you can either restore a backup of the queue manager or delete the queue manager.

If any transactions were active when the queue manager stopped, the local queues containing the persistent, uncommitted messages put or got inside these transactions are also required to start the queue manager successfully. If any of these local queues is found to be damaged, and the queue manager supports media recovery, it automatically tries to re-create them from their media images. If any of the queues cannot be recovered, IBM MQ cannot start.

If any damaged local queues containing uncommitted messages are discovered during startup processing on a queue manager that does not support media recovery, the queues are marked as damaged objects and the uncommitted messages on them are ignored. This situation is because it is not possible to perform media recovery of damaged objects on such a queue manager and the only action left is to delete them. Message AMQ7472 is issued to report any damage.

## Recovering damaged objects at other times

Media recovery of objects is automatic only during startup. At other times, when object damage is detected, operator message AMQ7472 is issued and most operations using the object fail. If the queue manager object is damaged at any time after the queue manager has started, the queue manager performs a pre-emptive shutdown. When an object has been damaged you can delete it or, if the queue manager is using a linear log, attempt to recover it from its media image using the **rcrmqobj** command (see **rcrmqobj** for further details).

## Protecting IBM MQ log files

Do not touch the log files when a queue manager is running, recovery might be impossible. Use super user or mqm authority to protect log files against inadvertent modification.

Do not remove the active log files manually when an IBM MQ queue manager is running. If a user inadvertently deletes the log files that a queue manager needs to restart, IBM MQ **does not** issue any errors and continues to process data *including persistent messages*. The queue manager shuts down normally, but can fail to restart. Recovery of messages then becomes impossible.

Users with the authority to remove logs that are being used by an active queue manager also have authority to delete other important queue manager resources (such as queue files, the object catalog, and IBM MQ executable files). They can therefore damage, perhaps through inexperience, a running or dormant queue manager in a way against which IBM MQ cannot protect itself.

Exercise caution when conferring super user or mqm authority.

## Dumping the contents of the log using the **dmpmqlog** command

How to use the **dmpmqlog** command to dump the contents of the queue manager log.

Use the **dmpmqlog** command to dump the contents of the queue manager log. By default all active log records are dumped, that is, the command starts dumping from the head of the log (usually the start of the last completed checkpoint).

The log can usually be dumped only when the queue manager is not running. Because the queue manager takes a checkpoint during shutdown, the active portion of the log usually contains a small number of log records. However, you can use the **dmpmqlog** command to dump more log records using one of the following options to change the start position of the dump:

- Start dumping from the *base* of the log. The base of the log is the first log record in the log file that contains the head of the log. The amount of additional data dumped in this case depends on where the head of the log is positioned in the log file. If it is near the start of the log file, only a small amount of additional data is dumped. If the head is near the end of the log file, significantly more data is dumped.
- Specify the start position of the dump as an individual log record. Each log record is identified by a unique *log sequence number (LSN)*. In the case of circular logging, this starting log record cannot be before the base of the log; this restriction does not apply to linear logs. You might need to reinstate inactive log files before running the command. You must specify a valid LSN, taken from previous **dmpmqlog** output, as the start position.

For example, with linear logging you can specify the `nextlsn` from your last **dmpmqlog** output. The `nextlsn` appears in Log File Header and indicates the LSN of the next log record to be written. Use this as a start position to format all log records written since the last time the log was dumped.

- **For linear logs only**, you can instruct **dmpmqlog** to start formatting log records from any given log file extent. In this case, **dmpmqlog** expects to find this log file, and each successive one, in the same directory as the active log files. This option does not apply to circular logs, where **dmpmqlog** cannot access log records prior to the base of the log.

The output from the **dmpmqlog** command is the Log File Header and a series of formatted log records. The queue manager uses several log records to record changes to its data.

Some of the information that is formatted is only of use internally. The following list includes the most useful log records:

### Log File Header

Each log has a single log file header, which is always the first thing formatted by the **dmpmqlog** command. It contains the following fields:

<i>logactive</i>	The number of primary log extents.
<i>loginactive</i>	The number of secondary log extents.
<i>logsize</i>	The number of 4 KB pages per extent.
<i>baselsn</i>	The first LSN in the log extent containing the head of the log.
<i>nextlsn</i>	The LSN of the next log record to be written.
<i>headlsn</i>	The LSN of the log record at the head of the log.
<i>tailsn</i>	The LSN identifying the tail position of the log.
<i>hflag1</i>	Whether the log is CIRCULAR or LOG RETAIN (linear).
<i>HeadExtentID</i>	The log extent containing the head of the log.

### Log Record Header

Each log record within the log has a fixed header containing the following information:

<i>LSN</i>	The log sequence number.
<i>LogRecdType</i>	The type of the log record.
<i>XTranid</i>	The transaction identifier associated with this log record (if any).  A <i>TranType</i> of MQI indicates an IBM MQ-only transaction. A <i>TranType</i> of XA is involved with other resource managers. Updates involved within the same unit of work have the same <i>XTranid</i> .
<i>QueueName</i>	The queue associated with this log record (if any).
<i>Qid</i>	The unique internal identifier for the queue.
<i>PrevLSN</i>	The LSN of the previous log record within the same transaction (if any).

### Start Queue Manager

This logs that the queue manager has started.

<i>StartDate</i>	The date that the queue manager started.
<i>StartTime</i>	The time that the queue manager started.

### Stop Queue Manager

This logs that the queue manager has stopped.

<i>StopDate</i>	The date that the queue manager stopped.
<i>StopTime</i>	The time that the queue manager stopped.
<i>ForceFlag</i>	The type of shutdown used.

### Start Checkpoint

This denotes the start of a queue manager checkpoint.

### End Checkpoint

This denotes the end of a queue manager checkpoint.

<i>ChkPtLSN</i>	The LSN of the log record that started this checkpoint.
-----------------	---------------------------------------------------------

### Put Message

This logs a persistent message put to a queue. If the message was put under sync point, the log record header contains a non-null *XTranid*. The remainder of the record contains:

<i>MapIndex</i>	An identifier for the message on the queue. It can be used to match the corresponding <b>MQGET</b> that was used to get this message from the queue. In this case a subsequent <i>Get Message</i> log record can be found containing the same <i>QueueName</i> and <i>MapIndex</i> . At this point the <i>MapIndex</i> identifier can be reused for a subsequent put message to that queue.
<i>Data</i>	Contained in the hex dump for this log record is various internal data followed by a representation of the Message Descriptor (eyecatcher MD) and then the message data itself.

### Put Part

Persistent messages that are too large for a single log record are logged as multiple *Put Part* log records followed by a single *Put Message* record. If there are *Put Part* records, then the *PrevLSN* field will chain the *Put Part* records and the final *Put Message* record together.

<i>Data</i>	Continues the message data where the previous log record left off.
-------------	--------------------------------------------------------------------

### Get Message

Only gets of persistent messages are logged. If the message was got under sync point, the log record header contains a non-null *XTranid*. The remainder of the record contains:

<i>MapIndex</i>	Identifies the message that was retrieved from the queue. The most recent <i>Put Message</i> log record containing the same <i>QueueName</i> and <i>MapIndex</i> identifies the message that was retrieved.
<i>QPriority</i>	The priority of the message retrieved from the queue.

### Start Transaction

Indicates the start of a new transaction. A *TranType* of MQI indicates an IBM MQ-only transaction. A *TranType* of XA indicates one that involves other resource managers. All updates made by this transaction will have the same *XTranid*.

### Prepare Transaction

Indicates that the queue manager is prepared to commit the updates associated with the specified *XTranid*. This log record is written as part of a two-phase commit involving other resource managers.

### Commit Transaction

Indicates that the queue manager has committed all updates made by a transaction.

### Rollback Transaction

This denotes the queue manager's intention to roll back a transaction.

### End Transaction

This denotes the end of a rolled-back transaction.

### Transaction Table

This record is written during sync point. It records the state of each transaction that has made persistent updates. For each transaction the following information is recorded:

<i>XTranid</i>	The transaction identifier.
<i>FirstLSN</i>	The LSN of the first log record associated with the transaction.
<i>LastLSN</i>	The LSN of the last log record associated with the transaction.

### Transaction Participants

This log record is written by the XA Transaction Manager component of the queue manager. It records the external resource managers that are participating in transactions. For each participant the following is recorded:

<i>RMName</i>	The name of the resource manager.
<i>RMID</i>	The resource manager identifier. This is also logged in subsequent <i>Transaction Prepared</i> log records that record global transactions in which the resource manager is participating.
<i>SwitchFile</i>	The switch load file for this resource manager.
<i>XAOpenString</i>	The XA open string for this resource manager.
<i>XACloseString</i>	The XA close string for this resource manager.

### Transaction Prepared

This log record is written by the XA Transaction Manager component of the queue manager. It indicates that the specified global transaction has been successfully prepared. Each of the participating resource managers will be instructed to commit. The *RMID* of each prepared resource manager is recorded in the log record. If the queue manager itself is participating in the transaction a *Participant Entry* with an *RMID* of zero will be present.

### Transaction Forget

This log record is written by the XA Transaction Manager component of the queue manager. It follows the *Transaction Prepared* log record when the commit decision has been delivered to each participant.

### Purge Queue

This logs the fact that all messages on a queue have been purged, for example, using the MQSC command CLEAR QUEUE.

### Queue Attributes

This logs the initialization or change of the attributes of a queue.

### Create Object

This logs the creation of an IBM MQ object.

<i>ObjName</i>	The name of the object that was created.
<i>UserId</i>	The user ID performing the creation.

### Delete Object

This logs the deletion of an IBM MQ object.

*ObjName*

The name of the object that was deleted.

## Backing up and restoring IBM MQ queue manager data

Backing up queue managers and queue manager data.

Periodically, you can take measures to protect queue managers against possible corruption caused by hardware failures. There are three ways of protecting a queue manager:

### Back up the queue manager data

If the hardware fails, a queue manager might be forced to stop. If any queue manager log data is lost due to the hardware failure, the queue manager might be unable to restart. If you back up queue manager data you might be able to recover some, or all, of the lost queue manager data.

In general, the more frequently you back up queue manager data, the less data you lose in the event of hardware failure that results in loss of integrity of the recovery log.

To back up queue manager data, the queue manager must not be running.

To back up and restore queue manager data see:

- “Backing up queue manager data” on page 1151.
- “Restoring queue manager data” on page 1151.

### Use a backup queue manager

If the hardware failure is severe, a queue manager might be unrecoverable. In this situation, if the unrecoverable queue manager has a dedicated backup queue manager, the backup queue manager can be activated in place of the unrecoverable queue manager. If it was updated regularly, the backup queue manager log can contain log data that includes the last complete log from the unrecoverable queue manager.

A backup queue manager can be updated while the existing queue manager is still running.

To create and activate a backup queue manager see:

- “Creating a backup queue manager” on page 1153.
- “Starting a backup queue manager” on page 1154.

### Back up the queue manager configuration only

If the hardware fails, a queue manager might be forced to stop. If both the queue manager configuration and log data is lost due to the hardware failure, the queue manager will be unable to restart or to be recovered from the log. If you back up the queue manager configuration you would be able to recreate the queue manager and all of its objects from saved definitions.

To back up queue manager configuration, the queue manager must be running.

To back up and restore the queue manager configuration see:

- “Backing up queue manager configuration” on page 1154
- “Restoring queue manager configuration” on page 1155

## Backing up queue manager data

Backing up queue manager data can help you to guard against possible loss of data caused by hardware errors.

### Before you begin

Ensure that the queue manager is not running. If you try to take a backup of a running queue manager, the backup might not be consistent because of updates in progress when the files were copied. If possible, stop your queue manager by running the **endmqm -w** command (a wait shutdown), only if that fails, use the **endmqm -i** command (an immediate shutdown).

### About this task

To take a backup copy of a queue manager's data, complete the following tasks:

1. Search for the directories under which the queue manager places its data and its log files, by using the information in the configuration files. For more information, see “Changing IBM MQ and queue manager configuration information” on page 755.

**Note:** The names that appear in the directory are transformed to ensure that they are compatible with the platform on which you are using IBM MQ. For more information about name transformations, see Understanding IBM MQ file names.

2. Take copies of all the queue manager's data and log file directories, including all subdirectories. Make sure that you do not miss any files, especially the log control file, as described in “What logs look like” on page 1130, and the configuration files as described in “Initialization and configuration files” on page 852. Some of the directories might be empty, but you need them all to restore the backup at a later date.
3. Preserve the ownerships of the files. For IBM MQ for UNIX and Linux systems, you can do this with the **tar** command. (If you have queues larger than 2 GB, you cannot use the **tar** command. For more information, see Enabling large queues.

**Note:** When you upgrade to IBM WebSphere MQ Version 7.5 and later, ensure to take a backup of the **.ini** file and the registry entries. The queue manager information is stored in the **.ini** file and can be used to revert to a previous version of IBM MQ.

## Restoring queue manager data

Follow these steps to restore a backup of a queue manager's data.

### Before you begin

Ensure that the queue manager is not running.

### About this task

To restore a backup of a queue manager's data:

1. Find the directories under which the queue manager places its data and its log files, by using the information in the configuration files.
2. Empty the directories into which you are going to place the backed-up data.
3. Copy the backed-up queue manager data and log files into the correct places.
4. Update the configuration information files.

Check the resulting directory structure to ensure that you have all the required directories.

For more information about IBM MQ directories and subdirectories, see Directory structure on Windows systems and Directory content on UNIX and Linux systems.

Make sure that you have a log control file as well as the log files. Also check that the IBM MQ and queue manager configuration files are consistent so that IBM MQ can look for the restored data in the correct places.

For circular logging, back up the queue manager data and log file directories at the same time so that you can restore a consistent set of queue manager data and logs.

For linear logging, back up the queue manager data and log file directories at the same time. It is possible to restore only the queue manager data files if a corresponding complete sequence of log files is available.

**Note:** When you upgrade to IBM WebSphere MQ Version 7.5 and later, ensure to take a backup of the `.ini` file and the registry entries. The queue manager information is stored in the `.ini` file and can be used to revert to a previous version of IBM MQ.

## Results

If the data was backed up and restored correctly, the queue manager will now start.

## Using a backup queue manager

An existing queue manager can have a dedicated backup queue manager.

A backup queue manager is an inactive copy of the existing queue manager. If the existing queue manager becomes unrecoverable due to severe hardware failure, the backup queue manager can be brought online to replace the unrecoverable queue manager.

The existing queue manager log files must regularly be copied to the backup queue manager to ensure that the backup queue manager remains an effective method for disaster recovery. The existing queue manager does not need to be stopped for log files to be copied, however you should only copy a log file if the queue manager has finished writing to it. Because the existing queue manager log is continually updated, there is always a slight discrepancy between the existing queue manager log and the log data copied to the backup queue manager log. Regular updates to the backup queue manager minimizes the discrepancy between the two logs.

If a backup queue manager is required to be brought online it must be activated, and then started. The requirement to activate a backup queue manager before it is started is a preventive measure to protect against a backup queue manager being started accidentally. After a backup queue manager is activated it can no longer be updated.

For information on how to create, update, and start a backup queue manager, see the following topics:

- “Creating a backup queue manager” on page 1153
- “Updating a backup queue manager” on page 1153
- “Starting a backup queue manager” on page 1154



## Creating a backup queue manager

You can only use a backup queue manager when using linear logging.

To create a backup queue manager for an existing queue manager, do the following:

1. Create a backup queue manager for the existing queue manager using the control command `strmqm`. The backup queue manager requires the following:
  - To have the same attributes as the existing queue manager, for example the queue manager name, the logging type, and the log file size.
  - To be on the same platform as the existing queue manager.
  - To be at an equal, or higher, code level than the existing queue manager.
2. Take copies of all the existing queue manager's data and log file directories, including all subdirectories, as described in "Backing up queue manager data" on page 1151.
3. Overwrite the backup queue manager's data and log file directories, including all subdirectories, with the copies taken from the existing queue manager.
4. Execute the following control command on the backup queue manager:

```
strmqm -r BackupQMName
```

This flags the queue manager as a backup queue manager within IBM MQ, and replays all the copied log extents to bring the backup queue manager in step with the existing queue manager.

## Updating a backup queue manager

To ensure that a backup queue manager remains an effective method for disaster recovery it must be updated regularly.

Regular updating lessens the discrepancy between the backup queue manager log, and the current queue manager log. There is no need to stop the queue manager to be backed up.

To update a backup queue manager, do the following:

1. Issue the following Script (MQSC) command on the queue manager to be backed up:

```
RESET QMGR TYPE(ADVANCELOG)
```

This stops any writing to the current log, and then advances the queue manager logging to the next log extent. This ensures you back up all information logged up to the current time.

2. Obtain the (new) current active log extent number by issuing the following Script (MQSC) command on the queue manager to be backed up:

```
DIS QMSTATUS CURRLOG
```

3. Copy the updated log extent files from the current queue manager log directory to the backup queue manager log directory - copy all the log extents since the last update, and up to (but not including) the current extent noted in step 2. Copy only log extent files, the ones beginning with "S...".
4. Issue the following control command on the backup queue manager:

```
strmqm -r BackupQMName
```

This replays all the copied log extents and brings the backup queue manager into step with the queue manager. When the replay finishes you receive a message that identifies all the log extents required for restart recovery, and all the log extents required for media recovery.

**Warning:** If you copy a **non-contiguous** set of logs to the backup queue manager log directory, only the logs up to the point where the first missing log is found will be replayed.

## Starting a backup queue manager

You can substitute a backup queue manager for an unrecoverable queue manager.

To do this, perform the following steps:

1. Execute the following control command to activate the backup queue manager:

```
strmqm -a BackupQMName
```

The backup queue manager is activated. Now active, the backup queue manager can no longer be updated.

2. Execute the following control command to start the backup queue manager:

```
strmqm BackupQMName
```

IBM MQ regards this as restart recovery, and utilizes the log from the backup queue manager. During the last update to the backup queue manager replay will have occurred, therefore only the active transactions from the last recorded checkpoint are rolled back.

When an unrecoverable queue manager is substituted for a backup queue manager some of the queue manager data from the unrecoverable queue manager can be lost. The amount of lost data is dependent on how recently the backup queue manager was last updated. The more recently the last update, the less queue manager data loss.

3. Restart all channels.

Check the resulting directory structure to ensure that you have all the required directories.

See Planning file system support for more information about IBM MQ directories and subdirectories.

Make sure that you have a log control file as well as the log files. Also check that the IBM MQ and queue manager configuration files are consistent so that IBM MQ can look in the correct places for the restored data.

If the data was backed up and restored correctly, the queue manager will now start.


**Note:** Even though the queue manager data and log files are held in different directories, back up and restore the directories at the same time. If the queue manager data and log files have different ages, the queue manager is not in a valid state and will probably not start. If it does start, your data is likely to be corrupt.

## Backing up queue manager configuration

Backing up queue manager configuration can help you to rebuild a queue manager from its definitions.

To take a backup copy of a queue manager's configuration:

1. Ensure that the queue manager is running.
2.
  - a. On AIX, HP-UX, Linux, Solaris, or Windows: Execute the Dump MQ Configuration command (dmpmqcfg) using the default formatting option of (-f mqsc) MQSC and all attributes (-a), use standard output redirection to store the definitions into a file, for example:




```
dmpmqcfg -m MYQMGR -a > /mq/backups/MYQMGR.mqsc
```
  - b.  On IBM i: Execute the Dump MQ Configuration command (DMPMQMCFG) using the default formatting option of OUTPUT(\*MQSC) and EXPATTR(\*ALL), use the TOFILE and TOMBR to store the definitions into a physical file member, for example:


```
DMPMQMCFG MQMNAME(MYQMGR) OUTPUT(*MQSC) EXPATTR(*ALL) TOFILE(QMQMSAMP/QMQSC) TOMBR(MYQMGRDEF)
```

## Restoring queue manager configuration

Follow these steps to restore a backup of a queue manager's configuration.

To restore a backup of a queue manager's configuration:

1. Ensure that the queue manager is running. Note that the queue manager may have been recreated if damage to the data and logs is unrecoverable by other means.
2. Depending on your platform, execute one of the following commands:
  - a.    On AIX, HP-UX, Linux, Solaris, or Windows: Execute **runmqsc** against the queue manager, use standard input redirection to restore the definitions from a script file generated by the Dump MQ Configuration (**dmpmqcfg**) command, for example:  

```
runmqsc MYQMGR < /mq/backups/MYQMGR.mqsc
```
  - b.  On IBM i: Execute **STRMQMMQSC** against the queue manager, use **SRCMBR** and **SRCFILE** to restore the definitions from the physical file member generated by the Dump MQ Configuration (**DMPMQMCFG**) command, for example:  

```
STRMQMMQSC MQMNAME(MYQMGR) SRCFILE(QMQMSAMP/QMQSC) SRCMBR(MYQMGR)
```

### Related information:

dmpmqcfg (dump queue manager configuration)

---

## Configuring JMS resources

One of the ways in which a JMS application can create and configure the resources that it needs to connect to IBM MQ and access destinations for sending or receiving messages is by using the Java Naming and Directory Interface (JNDI) to retrieve administered objects from a location within the naming and directory service that is called the JNDI namespace. Before a JMS application can retrieve administered objects from a JNDI namespace, you must first create and configure the administered objects.

### About this task

You can create and configure administered objects in IBM MQ by using either of the following tools:

#### MQ Explorer

You can use MQ Explorer to create and administer JMS object definitions that are stored in LDAP, in a local file system, or other locations.

#### IBM MQ JMS administration tool

The IBM MQ JMS administration tool is a command-line tool that you can use to create and configure IBM MQ JMS objects that are stored in LDAP, in a local file system, or other locations. The JMS administration tool uses a syntax that is similar to **runmqsc**, and also supports scripting.

The administration tool uses a configuration file to set the values of certain properties. A sample configuration file is supplied, which you can edit to suit your system before you start by using the tool to configure JMS resources. For more information about the configuration file, see “Configuring the JMS administration tool” on page 1160.

IBM MQ JMS applications that are deployed to WebSphere Application Server need to access JMS objects from the application server JNDI repository. Therefore, if you use JMS messaging between WebSphere Application Server and IBM MQ, you must create objects in WebSphere Application Server that correspond to the objects that you create in IBM MQ.

MQ Explorer and the IBM MQ JMS administration tool cannot be used to administer IBM MQ JMS objects that are stored in WebSphere Application Server. Instead, you can create and configure administered objects in WebSphere Application Server by using either of the following tools:

### **WebSphere Application Server administrative console**

The WebSphere Application Server administrative console is a web-based tool that you can use to manage IBM MQ JMS objects in WebSphere Application Server.

### **WebSphere Application Server wsadmin scripting client**

The WebSphere Application Server wsadmin scripting client provides specialized commands to administer IBM MQ JMS objects in WebSphere Application Server.

If you want to use a JMS application to access the resources of an IBM MQ queue manager from within WebSphere Application Server, you must use the IBM MQ messaging provider in WebSphere Application Server, which contains a version of the IBM MQ classes for JMS. The IBM MQ resource adapter that is supplied with WebSphere Application Server is used by all applications that carry out JMS messaging with the IBM MQ messaging provider. The IBM MQ resource adapter is usually updated automatically when you apply WebSphere Application Server fix packs, but if you have previously manually updated the resource adapter, you must manually update your configuration to ensure that maintenance is applied correctly.

#### **Related information:**

Writing IBM MQ classes for JMS applications

runmqsc

## **Configuring connection factories and destinations in a JNDI namespace**

JMS applications access administered objects in the naming and directory service through the Java Naming and Directory Interface (JNDI). The JMS administered objects are stored in a location within the naming and directory service that is referred to as the JNDI namespace. A JMS application can look up the administered objects to connect to IBM MQ and access destinations for sending or receiving messages.

### **About this task**

JMS applications look up the names of the JMS objects in the naming and directory service by using contexts:

#### **Initial context**

The initial context defines the root of the JNDI namespace. For each location in the naming and directory service, you need to specify an initial context to give a starting point from which a JMS application can resolve the names of the administered objects in that location of the naming and directory service.

#### **Subcontexts**

A context can have one or more subcontexts. A subcontext is a subdivision of a JNDI namespace and can contain administered objects such as connection factories and destinations as well as other subcontexts. A subcontext is not an object in its own right; it is merely an extension of the naming convention for the objects in the subcontext.

You can create contexts using either MQ Explorer or the IBM MQ JMS administration tool.

Before an IBM MQ classes for JMS application can retrieve administered objects from a JNDI namespace, you must first create the administered objects using either MQ Explorer or the IBM MQ JMS administration tool. You can create and configure the following types of JMS object:

#### **Connection factory**

A JMS connection factory object defines a set of standard configuration properties for connections. A JMS application uses a connection factory to create a connection to IBM MQ. You can create a connection factory that is specific to one of the two messaging domains, the point-to-point

messaging domain and the publish/subscribe messaging domain. Alternatively, from JMS 1.1, you can create domain-independent connection factories that can be used for both point-to-point and publish/subscribe messaging.

## Destination

A JMS destination is an object that represents the target of messages that the client produces and the source of messages that a JMS application consumes. The JMS application can either use a single destination object to put messages on and to get messages from, or the application can use separate destination objects. There are two types of destination object:

- JMS queue destination used in point-to-point messaging
- JMS topic destination used in publish/subscribe messaging

The following diagram shows an example of JMS objects created in an IBM MQ JNDI namespace.




Figure 162. JMS objects created in IBM MQ

If you use JMS messaging between WebSphere Application Server and IBM MQ, you must create corresponding objects in WebSphere Application Server to use to communicate with IBM MQ. When you create one of these objects in WebSphere Application Server, it is stored in the WebSphere Application Server JNDI namespace as shown in the following diagram.

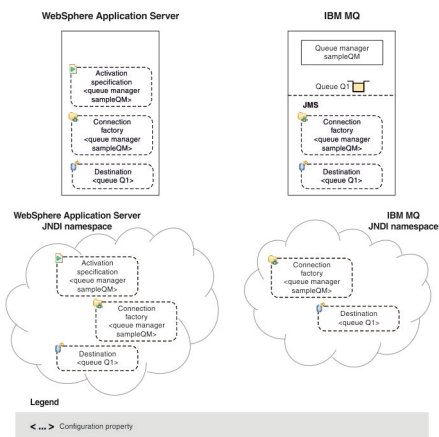


Figure 163. Objects created in WebSphere Application Server, and the corresponding objects in IBM MQ

If your application uses a message-driven bean (MDB), the connection factory is used for outbound messages only and inbound messages are received by an activation specification. Activation specifications are part of the Java EE Connector Architecture 1.5 (JCA 1.5) standard. JCA 1.5 provides a standard way to

integrate JMS providers, such as IBM MQ, with Java EE application servers such as WebSphere Application Server. A JMS activation specification can be associated with one or more message driven beans (MDBs) and provides the configuration necessary for these MDBs to listen for messages arriving at a destination.

You can use either the WebSphere Application Server administrative console or wsadmin scripting commands to create and configure the JMS resources that you need.

## Procedure

- To configure JMS objects for IBM MQ using MQ Explorer, see “Configuring JMS objects using MQ Explorer.”
- To configure JMS objects for IBM MQ using the IBM MQ JMS administration tool, see “Configuring JMS objects using the administration tool” on page 1159.
- To configure JMS objects for WebSphere Application Server, see “Configuring JMS resources in WebSphere Application Server” on page 1167.

## Results

An IBM MQ classes for JMS application can retrieve the administered objects from the JNDI namespace and, if required, set or change one or more of its properties by using either the IBM JMS extensions or the IBM MQ JMS extensions.

### Related information:

Using JNDI to retrieve administered objects in a JMS application

Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application

## Configuring JMS objects using MQ Explorer

Use the MQ Explorer graphical user interface to create JMS objects from IBM MQ objects, and IBM MQ objects from JMS objects, as well as for administering and monitoring other IBM MQ objects.

### About this task

MQ Explorer is the graphical user interface in which you can administer and monitor IBM MQ objects, whether they are hosted by your local computer or on a remote system. MQ Explorer runs on Windows and Linux x86-64. It can remotely connect to queue managers that are running on any supported platform including z/OS, enabling your entire messaging backbone to be viewed, explored, and altered from the console.

In MQ Explorer, all connection factories are stored in Connection Factories folders in the appropriate context and subcontexts.

You can perform the following types of task with MQ Explorer, either contextually from an existing object in the MQ Explorer, or from within a create new object wizard:

- Create a JMS Connection Factory from any of the following IBM MQ objects:
  - An IBM MQ queue manager, whether on your local computer or on a remote system.
  - An IBM MQ channel.
  - An IBM MQ listener.
- Add an IBM MQ queue manager to MQ Explorer using a JMS Connection Factory.
- Create a JMS queue from an IBM MQ queue.
- Create an IBM MQ queue from a JMS queue.
- Create a JMS topic from an IBM MQ topic, which can be an IBM MQ object or a dynamic topic.
- Create an IBM MQ topic from a JMS topic.

## Procedure

- Start MQ Explorer, if it is not already running. If MQ Explorer is running and displaying the Welcome page, close the Welcome page to start administering IBM MQ objects.
- If you have not already done so, create an initial context defining the root of the JNDI namespace in which the JMS objects are stored in the naming and directory service. When you have added the initial context to MQ Explorer, you can create connection factory objects, destination objects, and subcontexts in the JNDI namespace. The initial context is displayed in the Navigator view in the JMS Administered Objects folder. Note that although the full contents of the JNDI namespace are displayed, in MQ Explorer you can edit only the IBM MQ classes for JMS objects that are stored there. For more information, see Adding an initial context.
- Create and configure the subcontexts and JMS administered objects that you need. For more information, see Creating and configuring JMS administered objects.
- Configure IBM MQ. For more information, see Configuring IBM MQ using MQ Explorer .

### Related information:

Introduction to MQ Explorer

## Configuring JMS objects using the administration tool

You can use the IBM MQ JMS administration tool to define the properties of eight types of IBM MQ classes for JMS object and to store them within a JNDI namespace. Applications can then use JNDI to retrieve these administered objects from the namespace.

### About this task

The following table shows the eight types of administered objects that you can create, configure and manipulate using verbs. The Keyword column shows the strings that you can substitute for *TYPE* in the commands shown in Table 140.

Table 140. The JMS object types that are handled by the administration tool

Object Type	Keyword	Description
MQConnectionFactory	CF	The IBM MQ implementation of the JMS ConnectionFactory interface. This represents a factory object for creating connections in both the point-to-point and publish/subscribe domains.
MQQueueConnectionFactory	QCF	The IBM MQ implementation of the JMS QueueConnectionFactory interface. This represents a factory object for creating connections in the point-to-point domain.
MQTopicConnectionFactory	TCF	The IBM MQ implementation of the JMS TopicConnectionFactory interface. This represents a factory object for creating connections in the publish/subscribe domain.
MQQueue	Q	The IBM MQ implementation of the JMS Queue interface. This represents a destination for messages in the point-to-point domain.
MQTopic	T	The IBM MQ implementation of the JMS Topic interface. This represents a destination for messages in the publish/subscribe domain.
MQXAConnectionFactory <sup>1</sup>	XACF	The IBM MQ implementation of the JMS XAConnectionFactory interface. This represents a factory object for creating connections in both the point-to-point and publish/subscribe domains, and where the connections use the XA versions of JMS classes.

Table 140. The JMS object types that are handled by the administration tool (continued)

Object Type	Keyword	Description
MQXAQueueConnectionFactory <sup>1</sup>	XAQCF	The IBM MQ implementation of the JMS XAQueueConnectionFactory interface. This represents a factory object for creating connections in the point-to-point domain that use the XA versions of JMS classes.
MQXATopicConnectionFactory <sup>1</sup>	XATCF	The IBM MQ implementation of the JMS XATopicConnectionFactory interface. This represents a factory object for creating connections in the publish/subscribe domain that use the XA versions of JMS classes.

**Note:**

1. These classes are provided for use by vendors of application servers. They are unlikely to be directly useful to application programmers.

For more information about how to configure these objects, see “Configuring JMS objects” on page 1167.

The property types and values that you need to use this tool are listed in Properties of IBM MQ classes for JMS objects.

You can also use the tool to manipulate directory namespace subcontexts within the JNDI as described in “Configuring subcontexts” on page 1164.

You can also create and configure JMS administered objects with MQ Explorer.

**Related information:**

Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application

Using JNDI to retrieve administered objects in a JMS application

**Configuring the JMS administration tool**

The IBM MQ JMS administration tool uses a configuration file to set the values of certain properties. A sample configuration file is supplied, which you can edit to suit your system.

**About this task**

The configuration file is a plain-text file that consists of a set of key-value pairs, separated by the equal sign (=). You configure the administration tool by setting values for the three properties defined in the configuration file. The following example shows these three properties:

```
#Set the service provider
INITIAL_CONTEXT_FACTORY=com.sun.jndi.ldap.LdapCtxFactory
#Set the initial context
PROVIDER_URL=ldap://polaris/o=ibm_us,c=us
#Set the authentication type
SECURITY_AUTHENTICATION=none
```

(In this example, a hash sign (#) in the first column of the line indicates a comment, or a line that is not used.)

A sample configuration file, which is used as the default configuration file, is supplied with IBM MQ. The sample file is called JMSAdmin.config, and is found in the <MQ\_JAVA\_INSTALL\_PATH>/bin directory. You can either edit this sample file to define the settings needed for your system, or create your own configuration file.




When you start the administration tool, you can specify the configuration file that you want to use by using the `-cfg` command-line parameter, as described in “Starting the administration tool” on page 1162. If you do not specify a configuration file name when you invoke the tool, the tool attempts to load the default configuration file (`JMSAdmin.config`). It searches for this file first in the current directory, and then in the `<MQ_JAVA_INSTALL_PATH>/bin` directory, where `<MQ_JAVA_INSTALL_PATH>` is the path to your IBM MQ classes for JMS installation.

The names of JMS objects that are stored in an LDAP environment must comply with LDAP naming conventions. One of these conventions is that object and context names must include a prefix, such as `cn=` (common name), or `ou=` (organizational unit). The administration tool simplifies the use of LDAP service providers by allowing you to refer to object and context names without a prefix. If you do not supply a prefix, the tool automatically adds a default prefix to the name you supply. For LDAP, this is `cn=`. If required, you can change the default prefix by setting the **NAME\_PREFIX** property in the configuration file.

**Note:** You might need to configure your LDAP server to store Java objects. For more information, see the documentation for your LDAP server.

## Procedure

1. Define the service provider that the tool uses by configuring the **INITIAL\_CONTEXT\_FACTORY** property. The supported values for this property are as follows:
  - `com.sun.jndi.ldap.LdapCtxFactory` (for LDAP)
  - `com.sun.jndi.fscontext.RefFSContextFactory` (for file system context)
  -  `com.ibm.jndi.LDAPCtxFactory` is supported on z/OS only, and provides access to an LDAP server. However, this class is incompatible with `com.sun.jndi.ldap.LdapCtxFactory`, in that objects created using one `InitialContextFactory` cannot be read or modified using the other.

You can also use the administration tool to connect to other JNDI contexts by using three parameters defined in the `JMSAdmin` configuration file. To use a different `InitialContextFactory`:

- a. Set the **INITIAL\_CONTEXT\_FACTORY** property to the required class name.
- b. Define the behavior of the `InitialContextFactory` using the **USE\_INITIAL\_DIR\_CONTEXT**, **NAME\_PREFIX** and **NAME\_READABILITY\_MARKER** properties. The settings for these properties are described in the sample configuration file comments.

You do not need to define the **USE\_INITIAL\_DIR\_CONTEXT**, **NAME\_PREFIX** and **NAME\_READABILITY\_MARKER** properties if you use one of the supported **INITIAL\_CONTEXT\_FACTORY** values. However, you can give values to these properties if you want to override the system defaults. For example, if your objects are stored in an LDAP environment, you can change the default prefix that the tool adds to object and context names by setting the **NAME\_PREFIX** property to the required prefix.

If you omit one or more of the three `InitialContextFactory` properties, the administration tool provides suitable defaults based on the values of the other properties.

2. Define the URL of the initial context of the session by configuring the **PROVIDER\_URL** property. This URL is the root of all JNDI operations carried out by the tool. Two forms of this property are supported:
  - `ldap://hostname/contextname`
  - `file:[drive:]/pathname`

The format of the LDAP URL can vary, depending on your LDAP provider. See your LDAP documentation for more information.

3. Define whether JNDI passes security credentials to your service provider by configuring the **SECURITY\_AUTHENTICATION** property. This property is used only when an LDAP service provider is used and can take one of three values:

### **none (anonymous authentication)**

If you set this parameter to `none`, JNDI does not pass any security credentials to the service provider, and *anonymous authentication* is performed.

### simple (simple authentication)

If you set the parameter to `simple`, security credentials are passed through JNDI to the underlying service provider. These security credentials are in the form of a user distinguished name (User DN) and password.

### CRAM-MD5 (CRAM-MD5 authentication mechanism)

If you set the parameter to `CRAM-MD5`, security credentials are passed through JNDI to the underlying service provider. These security credentials are in the form of a user distinguished name (User DN) and password.

If you do not supply a valid value for the `SECURITY_AUTHENTICATION` property, the property defaults to `none`.

If security credentials are required, you are prompted for them when the tool initializes. You can avoid this by setting the `PROVIDER_USERDN` and `PROVIDER_PASSWORD` properties in the JMSAdmin configuration file.

**Note:** If you do not use these properties, the text typed, *including the password*, is echoed to the screen. This might have security implications.

The tool does no authentication itself; the authentication task is delegated to the LDAP server. The LDAP server administrator must set up and maintain access privileges to different parts of the directory. See your LDAP documentation for more information. If authentication fails, the tool displays an appropriate error message and terminates.

More detailed information about security and JNDI is in the documentation at Oracle's Java website ( Oracle Technology Network for Java Developers ).

## Starting the administration tool

The administration tool has a command-line interface that you can use either interactively, or to start a batch process.

### About this task

The interactive mode provides a command prompt where you can enter administration commands. In the batch mode, the command to start the tool includes the name of a file that contains an administration command script.

### Procedure

Interactive mode

- To start the tool in interactive mode, enter the following command:

```
JMSAdmin [-t] [-v] [-cfg config_filename]
```

where:

- t** Enables trace (default is trace off)

The trace file is generated in "%MQ\_JAVA\_DATA\_PATH%\errors ( Windows ) or /var/mqm/trace ( UNIX ). The name of the trace file is of the form:

```
mqjms_ PID.trc
```

where *PID* is the process ID of the JVM.

- v** Produces verbose output (default is terse output)

- cfg config\_filename**

Names an alternative configuration file. If this parameter is omitted, the default configuration file, `JMSAdmin.config`, is used. For more information about the configuration file, see "Configuring the JMS administration tool" on page 1160.

A command prompt is displayed, which indicates that the tool is ready to accept administration commands. This prompt initially appears as:

```
InitCtx>
```

indicating that the current context (that is, the JNDI context to which all naming and directory operations currently refer) is the initial context defined in the **PROVIDER\_URL** configuration parameter. For more information about this parameter, see “Configuring the JMS administration tool” on page 1160.

As you traverse the directory namespace, the prompt changes to reflect this, so that the prompt always displays the current context.

Batch mode

- To start the tool in batch mode, enter the following command:

```
JMSAdmin <test.scp
```

where `test.scp` is a script file that contains administration commands. For more information, see “Using administration commands.” The last command in the file must be the **END** command.

## Using administration commands

The administration tool accepts commands consisting of an administration verb and its appropriate parameters.

### About this task

The following table lists the administration verbs that you can use when entering commands with the administration tool.

*Table 141. Administration verbs*

Verb	Short form	Description
ALTER	ALT	Change at least one of the properties of an administered object
DEFINE	DEF	Create and store an administered object, or create a subcontext
DISPLAY	DIS	Display the properties of one or more stored administered objects, or the contents of the current context
DELETE	DEL	Remove one or more administered objects from the namespace, or remove an empty subcontext
CHANGE	CHG	Alter the current context, allowing the user to traverse the directory namespace anywhere below the initial context (pending security clearance)
COPY	CP	Make a copy of a stored administered object, storing it under an alternative name
MOVE	MV	Alter the name under which an administered object is stored
END		Close the administration tool

### Procedure

- If the administration tool is not already started, start it as described in “Starting the administration tool” on page 1162. The command prompt is displayed, indicating that the tool is ready to accept administration commands. This prompt initially appears as:

```
InitCtx>
```

To change the current context, use the **CHANGE** verb as described in “Configuring subcontexts” on page 1164.

- Enter commands in the following form:

```
verb [param]*
```

where **verb** is one of the administration verbs listed in Table 141 on page 1163. All valid commands contain one verb, which appears at the beginning of the command in either its standard or short form. Verb names are not case-sensitive.

- To terminate a command, press Enter, unless you want to enter several commands together, in which case type the plus sign (+) directly before pressing Enter. Typically, to terminate commands, you press Enter. However, you can override this by typing the plus sign (+) directly before pressing Enter. This enables you to enter multiline commands, as shown in the following example:

```
DEFINE Q(BookingsInputQueue) +
QMGR(QM.POLARIS.TEST) +
QUEUE(BOOKINGS.INPUT.QUEUE) +
PORT(1415) +
CCSID(437)
```

- To close the administration tool, use the **END** verb. This verb cannot take any parameters.

## Configuring subcontexts

You can use the verbs **CHANGE**, **DEFINE**, **DISPLAY** and **DELETE** to configure directory namespace subcontexts.

### About this task

The use of these verbs is described in the following table.

Table 142. Syntax and description of commands used to manipulate subcontexts

Command syntax	Description
DEFINE CTX(ctxName)	Attempts to create a child subcontext of the current context, having the name ctName. Fails if there is a security violation, if the subcontext already exists, or if the name supplied is not valid.
DISPLAY CTX	Displays the contents of the current context. Administered objects are annotated with a, subcontexts with [D]. The Java type of each object is also displayed.
DELETE CTX(ctxName)	Attempts to delete the current context's child context having the name ctName. Fails if the context is not found, is non-empty, or if there is a security violation.
CHANGE CTX(ctxName)	<p>Alters the current context, so that it now refers to the child context having the name ctName. One of two special values of ctName can be supplied:</p> <p><b>=UP</b> moves to the parent of the current context</p> <p><b>=INIT</b> moves directly to the initial context</p> <p>Fails if the specified context does not exist, or if there is a security violation.</p>

The names of JMS objects that are stored in an LDAP environment must comply with LDAP naming conventions. One of these conventions is that object and context names must include a prefix, such as **cn=** (common name), or **ou=** (organizational unit). The administration tool simplifies the use of LDAP service providers by allowing you to refer to object and context names without a prefix. If you do not supply a prefix, the tool automatically adds a default prefix to the name you supply. For LDAP, this is **cn=**. If required, you can change the default prefix by setting the **NAME\_PREFIX** property in the configuration file. For more information, see “Configuring the JMS administration tool” on page 1160.

**Note:** You might need to configure your LDAP server to store Java objects. For more information, see the documentation for your LDAP server.

## Creating JMS objects

To create JMS connection factory and destination objects and store them in a JNDI namespace, use the `DEFINE` verb. To store your objects in an LDAP environment, you must give them names that comply with certain conventions. The administration tool can help you obey LDAP naming conventions by adding a default prefix to object names.

### About this task

The `DEFINE` verb creates an administered object with the type, name and properties that you specify. The new object is stored in the current context.

The names of JMS objects that are stored in an LDAP environment must comply with LDAP naming conventions. One of these conventions is that object and context names must include a prefix, such as `cn=` (common name), or `ou=` (organizational unit). The administration tool simplifies the use of LDAP service providers by allowing you to refer to object and context names without a prefix. If you do not supply a prefix, the tool automatically adds a default prefix to the name you supply. For LDAP, this is `cn=`. If required, you can change the default prefix by setting the `NAME_PREFIX` property in the configuration file. For more information, see “Configuring the JMS administration tool” on page 1160.

**Note:** You might need to configure your LDAP server to store Java objects. For more information, see the documentation for your LDAP server.

### Procedure

1. If the administration tool is not already started, start it as described in “Starting the administration tool” on page 1162. The command prompt is displayed, indicating that the tool is ready to accept administration commands.
2. Make sure that command prompt is showing the context in which you want to create the new object. When you start the administration tool, the prompt initially appears as:

```
InitCtx>
```

To change the current context, use the `CHANGE` verb as described in “Configuring subcontexts” on page 1164.

3. To create a connection factory, queue destination or topic destination, use the following command syntax:

```
DEFINE TYPE (name) [property]*
```

That is, type the `DEFINE` verb, followed by a `TYPE (name)` administered object reference, followed by zero or more *properties* (see Properties of IBM MQ classes for JMS objects ).

4. To create a connection factory, queue destination or topic destination, use the following command syntax:

```
DEFINE TYPE (name) [property]*
```

5. To display the newly created object, use the `DISPLAY` verb with the following command syntax:

```
DISPLAY TYPE (name)
```

### Example

The following example shows a queue called `testQueue` created in the initial context using the `DEFINE` verb. Since this object is being stored in an LDAP environment, although the object name `testQueue` is not entered with a prefix, the tool automatically adds one to ensure compliance with the LDAP naming convention. Submitting the command `DISPLAY Q(testQueue)` also causes this prefix to be added.

```
InitCtx> DEFINE Q(testQueue)
```

```
InitCtx> DISPLAY CTX
```

```
Contents of InitCtx
```

```
a cn=testQueue com.ibm.mq.jms.MQQueue
```

```
1 Object(s)
0 Context(s)
1 Binding(s), 1 Administered
```

### Sample error conditions creating a JMS object:

A number of common error conditions can arise when you create an object.

The following are examples of these error conditions:

#### CipherSpec mapped to CipherSuite

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) SSLCIPHERSUITE(RC4_MD5_US)
WARNING: Converting CipherSpec RC4_MD5_US to
CipherSuite SSL_RSA_WITH_RC4_128_MD5
```

#### Invalid property for object

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) PRIORITY(4)
Unable to create a valid object, please check the parameters supplied
Invalid property for a QCF: PRI
```

#### Invalid type for property value

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) CCSID(english)
Unable to create a valid object, please check the parameters supplied
Invalid value for CCS property: English
```

#### Property clash - client/bindings

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) HOSTNAME(polaris.hursley.ibm.com)
Unable to create a valid object, please check the parameters supplied
Invalid property in this context: Client-bindings attribute clash
```

#### Property clash - Exit initialization

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) SECEXITINIT(initStr)
Unable to create a valid object, please check the parameters supplied
Invalid property in this context: ExitInit string supplied
without Exit string
```

#### Property value outside valid range

```
InitCtx/cn=Trash> DEFINE Q(testQ) PRIORITY(12)
Unable to create a valid object, please check the parameters supplied
Invalid value for PRI property: 12
```

#### Unknown property

```
InitCtx/cn=Trash> DEFINE QCF(testQCF) PIZZA(ham and mushroom)
Unable to create a valid object, please check the parameters supplied
Unknown property: PIZZA
```

The following are examples of error conditions that might arise on Windows when looking up JNDI administered objects from a JMS application.

1. If you are using the WebSphere JNDI provider, `com.ibm.websphere.naming.WsnInitialContextFactory`, you must use a forward slash (/) to access administered objects defined in subcontexts; for example, `jms/MyQueueName`. If you use a backslash (\), an `InvalidNameException` is thrown.
2. If you are using the Oracle JNDI provider, `com.sun.jndi.fscontext.RefFSContextFactory`, you must use a backslash (\) to access administered objects defined in subcontexts; for example, `ctx1\\fred`. If you use a forward slash (/), a `NameNotFoundException` is thrown.

## Configuring JMS objects

You can use the verbs ALTER, DEFINE, DISPLAY, DELETE, COPY, and MOVE to manipulate administered objects in the directory namespace.

### About this task

Table 143 summarizes the use of these verbs. Substitute *TYPE* with the keyword that represents the required administered object, as described in “Configuring JMS objects using the administration tool” on page 1159.

Table 143. Syntax and description of commands used to manipulate administered objects

Command syntax	Description
ALTER <i>TYPE</i> (name) [property]*	Attempts to update the properties of the administered object with the ones supplied. Fails if there is a security violation, if the specified object cannot be found, or if the new properties supplied are not valid.
DEFINE <i>TYPE</i> (name) [property]*	Attempts to create an administered object of type <i>TYPE</i> with the supplied properties, and store it under the name <i>name</i> in the current context. Fails if there is a security violation, if the supplied name is not valid or an object of that name exists, or if the properties supplied are not valid.
DISPLAY <i>TYPE</i> (name)	Displays the properties of the administered object of type <i>TYPE</i> , bound under the name <i>name</i> in the current context. Fails if the object does not exist, or if there is a security violation.
DELETE <i>TYPE</i> (name)	Attempts to remove the administered object of type <i>TYPE</i> , having the name <i>name</i> , from the current context. Fails if the object does not exist, or if there is a security violation.
COPY <i>TYPE</i> (nameA) <i>TYPE</i> (nameB)	Makes a copy of the administered object of type <i>TYPE</i> , having the name <i>nameA</i> , naming the copy <i>nameB</i> . This all occurs within the scope of the current context. Fails if the object to be copied does not exist, if an object of name <i>nameB</i> exists, or if there is a security violation.
MOVE <i>TYPE</i> (nameA) <i>TYPE</i> (nameB)	Moves (renames) the administered object of type <i>TYPE</i> , having the name <i>nameA</i> , to <i>nameB</i> . This all occurs within the scope of the current context. Fails if the object to be moved does not exist, if an object of name <i>nameB</i> exists, or if there is a security violation.

## Configuring JMS resources in WebSphere Application Server

To configure JMS resources in WebSphere Application Server, you can either use the administrative console or wsadmin commands.

### About this task

Java Message Service (JMS) applications typically rely on externally configured objects which describe how the application connects to its JMS provider and the destinations it accesses. JMS applications use the Java Naming and Directory Interface (JNDI) to access the following types of object at runtime:

- Activation specifications (used by Java EE application servers)
- Unified connection factories (with JMS 1.1, domain-independent (unified) connection factories are preferred to domain-specific queue connection factories and topic connection factories)
- Topic connection factories (used by JMS 1.0 applications)
- Queue connection factories (used by JMS 1.0 applications)
- Queues
- Topics

Through the IBM MQ messaging provider in WebSphere Application Server, Java Message Service (JMS) messaging applications can use your IBM MQ system as an external provider of JMS messaging resources. To enable this approach, you configure the IBM MQ messaging provider in WebSphere Application Server to define JMS resources for connecting to any queue manager on the IBM MQ network.

You can use WebSphere Application Server to configure IBM MQ resources for applications (for example queue connection factories) and to manage messages and subscriptions associated with JMS destinations. You administer security through IBM MQ.

#### **Related information for WebSphere Application Server Version 8.5.5**

Interoperation using the IBM MQ messaging provider

Managing messaging with the IBM MQ messaging provider

Mapping of administrative console panel names to command names and IBM MQ names

#### **Related information for WebSphere Application Server Version 8.0**

Interoperation using the IBM MQ messaging provider

Managing messaging with the IBM MQ messaging provider

Mapping of administrative console panel names to command names and IBM MQ names

#### **Related information for WebSphere Application Server Version 7.0**

Interoperation using the IBM MQ messaging provider

Managing messaging with the IBM MQ messaging provider

Mapping of administrative console panel names to command names and IBM MQ names

## **Configuring JMS resources using the administrative console**

You can use the WebSphere Application Server administrative console to configure activation specifications, connection factories and destinations for the IBM MQ JMS provider.

### **About this task**

You can use the WebSphere Application Server administrative console to create, view, or modify any of the following resources:

- Activation specifications
- Domain-independent connection factories (JMS 1.1 or later)
- Queue connection factories
- Topic connection factories
- Queues
- Topics

The following steps provide an overview of the ways in which you can use the administrative console to configure JMS resources for use with the IBM MQ messaging provider. Each step includes the name of the topic in the WebSphere Application Server product documentation to which you can refer for more information. See *Related links* for links to these topics in the WebSphere Application Server Version 8.5.5, Version 8.0 and Version 7.0 product documentation.

In a mixed-version WebSphere Application Server cell, you can administer IBM MQ resources on nodes of all versions. However, some properties are not available on all versions. In this situation, only the properties of that particular node are displayed in the administrative console.

### **Procedure**

To create or configure an activation specification for use with the IBM MQ messaging provider:

- To create an activation specification, use the Create IBM MQ JMS resource wizard. You can either use the wizard to specify all the details for the activation specification, or you can choose to specify the



connection details for the IBM MQ by using a client channel definition table (CCDT). When you specify the connection details using the wizard, you can choose either to enter host and port information separately or, if you are using a multi-instance queue manager, to enter host and port information in the form of a connection name list. For more information, see *Creating an activation specification for the IBM MQ messaging provider*.

- To view or change the configuration properties of an activation specification, use the administrative console IBM MQ messaging provider connection factory settings panel. These configuration properties control how connections are created to associated queues and topics. For more information, see *Configuring an activation specification for the IBM MQ messaging provider*.

To create or configure a unified connection factory, a queue connection factory, or a topic connection factory for use with the IBM MQ messaging provider:

- To create a connection factory, first select the type of connection factory that you want to create, then use the Create IBM MQ JMS resource wizard to specify the details.
  - If your JMS application is intended to use only point-to-point messaging, create a domain-specific connection factory for the point-to-point messaging domain that can be used for creating connections specifically for point-to-point messaging.
  - If your JMS application is intended only to use publish/subscribe messaging, create a domain-specific connection factory for the publish/subscribe messaging domain that can be used for creating connections specifically for publish/subscribe messaging.
  - For JMS 1.1 or later, create a domain-independent connection factory that can be used for both point-to-point messaging and publish/subscribe messaging, allowing your application to perform both point-to-point and publish/subscribe work under the same transaction.

You can choose whether to use the wizard to specify all the details for the connection factory, or you can choose to specify the connection details for the IBM MQ by using a client channel definition table (CCDT). When you specify the connection details using the wizard, you can choose either to enter host and port information separately or, if you are using a multi-instance queue manager, to enter host and port information in the form of a connection name list. For more information, see *Creating a connection factory for the IBM MQ messaging provider*.

To view or change the configuration properties of a connection factory:

- Use the administrative console connection factory settings panel for the type of connection factory that you want to configure. The configuration properties control how connections are created to associated queues and topics. For more information, see *Configuring a collection factory for the IBM MQ messaging provider*, or *Configuring a queue collection factory for the IBM MQ messaging provider*, or *Configuring a topic collection factory for the IBM MQ messaging provider*.

To configure a JMS queue destination for point-to-point messaging with the IBM MQ messaging provider:

- Use the administrative console IBM MQ messaging provider queue settings panel to define the following types of property:
  - General properties, including administration and IBM MQ queue properties.
  - Connection properties that specify how to connect to the queue manager that hosts the queue.
  - Advanced properties that control the behavior of connections made to IBM MQ messaging provider destinations.
  - Any custom properties for the queue destination.

For more information, see *Configuring a queue for the IBM MQ messaging provider*.

To create or configure a JMS topic destination for publish/subscribe messaging with the IBM MQ messaging provider:

- Use the IBM MQ messaging provider topic settings panel to define the following types of property:
  - General properties, including administration and IBM MQ topic properties.
  - Advanced properties that control the behavior of connections made to IBM MQ messaging provider destinations.
  - Any custom properties for the queue destination.

For more information, see *Configuring a topic for the IBM MQ messaging provider*.

**Related concepts:**

“Client channel definition table” on page 716

The client channel definition table (CCDT) determines the channel definitions and authentication information used by client applications to connect to the queue manager. On platforms other than z/OS a CCDT is created automatically. You must then make it available to the client application.

“Multi-instance queue managers” on page 1075

Multi-instance queue managers are instances of the same queue manager configured on different servers. One instance of the queue manager is defined as the active instance and another instance is defined as the standby instance. If the active instance fails, the multi-instance queue manager restarts automatically on the standby server.

“Configuring publish/subscribe messaging” on page 1004

You can start, stop and display the status of queued publish/subscribe. You can also add and remove streams, and add and delete queue managers from a broker hierarchy.

**Related information for WebSphere Application Server Version 8.5.5**

IBM MQ messaging provider activation specifications

Creating an activation specification for the IBM MQ messaging provider

Configuring an activation specification for the IBM MQ messaging provider

Creating a connection factory for the IBM MQ messaging provider

Configuring a unified connection factory for the IBM MQ messaging provider

Configuring a queue connection factory for the IBM MQ messaging provider

Configuring a topic connection factory for the IBM MQ messaging provider

Configuring a queue for the IBM MQ messaging provider

Configuring a topic for the IBM MQ messaging provider

**Related information for WebSphere Application Server Version 8.0**

IBM MQ messaging provider activation specifications

Creating an activation specification for the IBM MQ messaging provider

Configuring an activation specification for the IBM MQ messaging provider

Creating a connection factory for the IBM MQ messaging provider

Configuring a unified connection factory for the IBM MQ messaging provider

Configuring a queue connection factory for the IBM MQ messaging provider

Configuring a topic connection factory for the IBM MQ messaging provider

Configuring a queue for the IBM MQ messaging provider

Configuring a topic for the IBM MQ messaging provider

**Related information for WebSphere Application Server Version 7.0**

IBM MQ messaging provider activation specifications

Creating an activation specification for the IBM MQ messaging provider

Configuring an activation specification for the IBM MQ messaging provider

Creating a connection factory for the IBM MQ messaging provider

Configuring a unified connection factory for the IBM MQ messaging provider

Configuring a queue connection factory for the IBM MQ messaging provider

Configuring a topic connection factory for the IBM MQ messaging provider

Configuring a queue for the IBM MQ messaging provider

Configuring a topic for the IBM MQ messaging provider

**Configuring JMS resources using wsadmin scripting commands**

You can use WebSphere Application Server wsadmin scripting commands to create, modify, delete or show information about JMS activation specifications, connection factories, queues and topics. You can also display and manage the settings for the IBM MQ resource adapter.

## About this task

The following steps provide an overview of the ways in which you can use WebSphere Application Server wsadmin commands to configure JMS resources for use with the IBM MQ messaging provider. For more information about how to use these commands, see *Related links* for links to the WebSphere Application Server Version 8.5.5, Version 8.0 and Version 7.0 product documentation.

To run a command, use the AdminTask object of the wsadmin scripting client.

After using a command to create a new object or make changes, save your changes to the master configuration. For example, use the following command:

```
AdminConfig.save()
```

To see a list of the available IBM MQ messaging provider administrative commands, plus a brief description of each command, enter the following command at the wsadmin prompt:

```
print AdminTask.help('WMQAdminCommands')
```

To see overview help on a given command, enter the following command at the wsadmin prompt:

```
print AdminTask.help('command_name')
```

## Procedure

To list all of the IBM MQ messaging provider resources defined at the scope at which a command is issued, use the following commands.

- To list the activation specifications, use the **listWMQActivationSpecs** command.
- To list the connection factories, use the **listWMQConnectionFactory** command.
- To list the queue type destinations, use the **listWMQQueues** command.
- To list the topic type destinations, use the **listWMQTopics** command.

To create a JMS resource for the IBM MQ messaging provider at a specific scope, use the following commands.

- To create an activation specification, use the **createWMQActivationSpec** command. You can either create an activation specification by specifying all the parameters to be used for establishing a connection, or you can create the activation specification so that it uses a client channel definition table (CCDT) to locate the queue manager to connect to.
- To create a connection factory, use the **createWMQConnectionFactory** command, using the **-type** parameter to specify the type of connection factory that you want to create:
  - If your JMS application is intended to use only point-to-point messaging, create a domain-specific connection factory for the point-to-point messaging domain that can be used for creating connections specifically for point-to-point messaging.
  - If your JMS application is intended only to use publish/subscribe messaging, create a domain-specific connection factory for the publish/subscribe messaging domain that can be used for creating connections specifically for publish/subscribe messaging.
  - For JMS 1.1 or later, create a domain-independent connection factory that can be used for both point-to-point messaging and publish/subscribe messaging, allowing your application to perform both point-to-point and publish/subscribe work under the same transaction.

The default type is domain-independent connection factory.

- To create a queue type destination, use the **createWMQQueue** command.
- To create a topic type destination, use the **createWMQTopic** command.

To modify a JMS resource for the IBM MQ messaging provider at a specific scope, use the following commands.

- To modify an activation specification, use the **modifyWMQActivationSpec** command. You cannot change the type of an activation specification. For example, you cannot create an activation specification where you enter all the configuration information manually and then modify it to use a CCDT.
- To modify a connection factory, use the **modifyWMQConnectionFactory** command.
- To modify a queue type destination, use the **modifyWMQQueue** command.
- To modify a topic type destination, use the **modifyWMQTopic** command.

To delete a JMS resource for the IBM MQ messaging provider at a specific scope, use the following commands.

- To delete an activation specification, use the **deleteWMQActivationSpec** command.
- To delete a connection factory, use the **deleteWMQConnectionFactory** command.
- To delete a queue type destination, use the **deleteWMQQueue** command.
- To delete a topic type destination, use the **deleteWMQTopic** command.

To display information about a specific IBM MQ messaging provider resource, use the following commands.

- To display all the parameters, and their values, associated with a particular activation specification, use the **showWMQActivationSpec** command.
- To display all the parameters, and their values, associated with a particular connection factory, use the **showWMQConnectionFactory** command.
- To display all the parameters, and their values, associated with a particular queue type destination, use the **showWMQQueue** command.
- To display all the parameters, and their values, associated with a topic type destination, use the **showWMQTopic** command.

To manage settings for the IBM MQ resource adapter or the IBM MQ messaging provider, use the following commands.

- To manage the settings of the IBM MQ resource adapter that is installed at a particular scope, use the **manageWMQ** command.
- To display all the parameters, and their values that can be set by the **manageWMQ** command, use the **showWMQ** command. These settings are either related to the IBM MQ resource adapter or the IBM MQ messaging provider. The **showWMQ** command also shows any custom properties that are set on the IBM MQ resource adapter.

#### Related concepts:

“Client channel definition table” on page 716

The client channel definition table (CCDT) determines the channel definitions and authentication information used by client applications to connect to the queue manager. On platforms other than z/OS a CCDT is created automatically. You must then make it available to the client application.

“Multi-instance queue managers” on page 1075

Multi-instance queue managers are instances of the same queue manager configured on different servers. One instance of the queue manager is defined as the active instance and another instance is defined as the standby instance. If the active instance fails, the multi-instance queue manager restarts automatically on the standby server.

“Configuring publish/subscribe messaging” on page 1004

You can start, stop and display the status of queued publish/subscribe. You can also add and remove streams, and add and delete queue managers from a broker hierarchy.

#### Related information for WebSphere Application Server Version 8.5.5

- createWMQActivationSpec** command
- createWMQConnectionFactory** command
- createWMQQueue** command
- createWMQTopic** command
- deleteWMQActivationSpec** command
- deleteWMQConnectionFactory** command

**deleteWMQQueue** command  
**deleteWMQTopic** command  
**listWMQActivationSpecs** command  
**listWMQConnectionFactory** command  
**listWMQQueues** command  
**listWMQTopics** command  
**modifyWMQActivationSpec** command  
**modifyWMQConnectionFactory** command  
**modifyWMQQueue** command  
**modifyWMQTopic** command  
**showWMQActivationSpec** command  
**showWMQConnectionFactory** command  
**showWMQQueue** command  
**showWMQTopic** command  
**showWMQ** command  
**manageWMQ** command

#### **Related information for WebSphere Application Server Version 8.0**

**createWMQActivationSpec** command  
**createWMQConnectionFactory** command  
**createWMQQueue** command  
**createWMQTopic** command  
**deleteWMQActivationSpec** command  
**deleteWMQConnectionFactory** command  
**deleteWMQQueue** command  
**deleteWMQTopic** command  
**listWMQActivationSpecs** command  
**listWMQConnectionFactory** command  
**listWMQQueues** command  
**listWMQTopics** command  
**modifyWMQActivationSpec** command  
**modifyWMQConnectionFactory** command  
**modifyWMQQueue** command  
**modifyWMQTopic** command  
**showWMQActivationSpec** command  
**showWMQConnectionFactory** command  
**showWMQQueue** command  
**showWMQTopic** command  
**showWMQ** command  
**manageWMQ** command

#### **Related information for WebSphere Application Server Version 7.0**

**createWMQActivationSpec** command  
**createWMQConnectionFactory** command  
**createWMQQueue** command  
**createWMQTopic** command  
**deleteWMQActivationSpec** command

`deleteWMQConnectionFactory` command  
`deleteWMQQueue` command  
`deleteWMQTopic` command  
`listWMQActivationSpecs` command  
`listWMQConnectionFactories` command  
`listWMQQueues` command  
`listWMQTopics` command  
`modifyWMQActivationSpec` command  
`modifyWMQConnectionFactory` command  
`modifyWMQQueue` command  
`modifyWMQTopic` command  
`showWMQActivationSpec` command  
`showWMQConnectionFactory` command  
`showWMQQueue` command  
`showWMQTopic` command  
`showWMQ` command  
`manageWMQ` command

## Configuring the application server to use the latest resource adapter maintenance level

To ensure that the IBM MQ resource adapter is automatically updated to the latest available maintenance level when you apply WebSphere Application Server fix packs, you can configure all servers in your environment to use the latest version of the resource adapter contained in the WebSphere Application Server fix pack that you have applied to the installation of each node.

### Before you begin

**Important:** If you are using WebSphere Application Server Version 7.0, Version 8 or Version 8.5 on any platform, do not install the IBM MQ Version 8.0 resource adapter into the application server. The IBM MQ Version 8.0 resource adapter can only be deployed into an application server that supports JMS 2.0. However, WebSphere Application Server Version 7.0, Version 8 and Version 8.5 only support JMS 1.1. These versions of WebSphere Application Server come with the IBM WebSphere MQ Version 7.0 resource adapter, which can be used to connect to a IBM MQ Version 8.0 queue manager using either the BINDINGS or CLIENT transport.

### About this task

Use this task if any of the following circumstances apply to your configuration, and you want to configure all servers in your environment to use the latest version of the IBM MQ resource adapter:

- The JVM logs of any application server in your environment show the following IBM MQ resource adapter version information after WebSphere Application Server Version 7.0 Fix Pack 1 or later has been applied:  
WMSG1703I:RAR implementation Version 7.0.0.0-k700-L080820
- The JVM logs of any application server in your environment contain the following entry:  
WMSG1625E: It was not possible to detect  
the IBM MQ messaging provider code at the specified path <null>
- One or more nodes has previously been manually updated to use a specific maintenance level of the IBM MQ resource adapter that is now superseded by the latest version of the resource adapter contained in the current WebSphere Application Server maintenance level.

The *profile\_root* directory that the examples refer to is the home directory for the WebSphere Application Server profile, for example C:\Program Files\IBM\WebSphere\AppServer1.

When you have performed the following steps for all cells and single server installations in your environment, your servers automatically receive maintenance to the IBM MQ resource adapter when a new WebSphere Application Server fix pack is applied.

## Procedure

1. Start the application server. If the profile is part of a network deployment configuration, start the deployment manager and all node agents. If the profile contains an administrative agent, start the administrative agent.
2. Check the maintenance level of the IBM MQ resource adapter.
  - a. Open a command prompt window and change to the *profile\_root*\bin directory. For example, enter `cd C:\Program Files\IBM\WebSphere\AppServer1\bin`.
  - b. Start the wsadmin tool by entering `wsadmin.bat -lang jython`, then if prompted to do so, enter your username and password.
  - c. Type the following command, then press Return twice:

```
wmqInfoMBeansUnsplit = AdminControl.queryNames("WebSphere:type=WMQInfo,*")
wmqInfoMBeansSplit = AdminUtilities.convertToList(wmqInfoMBeansUnsplit)
for wmqInfoMBean in wmqInfoMBeansSplit: print wmqInfoMBean; print AdminControl.invoke(wmqInfoMBean, 'getInfo', '')
```

You can also run this command in Jacl. For further information about how to do this, see *Ensuring that servers use the latest available IBM MQ resource adapter maintenance level* in the WebSphere Application Server product documentation.

- d. Find the WMSG1703I message in the displayed output from the command and check the resource adapter level. For example, for WebSphere Application Server Version 7.0 Fixpack 15, the message should be:

```
WMSG1703I: RAR implementation Version 7.0.1.3-k701-103-100812
```

This message shows that the version is 7.0.1.3-k701-103-100812, which is the correct resource adapter level for this fixpack. However, if the following message is displayed instead, this means that you need to adjust the resource adapter to the correct level of maintenance for Fix Pack 15.

```
WMSG1703I: RAR implementation Version 7.0.0.0-k700-L080820
```

3. Copy the following Jython script into a file called `convertWMQRA.py`, then save it into the profile root directory, for example C:\Program Files\IBM\WebSphere\AppServer1\bin.

```
ras = AdminUtilities.convertToList(AdminConfig.list('J2CResourceAdapter'))

for ra in ras :
 desc = AdminConfig.showAttribute(ra, "description")
 if (desc == "WAS 7.0 Built In IBM MQ Resource Adapter") or (desc == "WAS 7.0.0.1 Built In IBM MQ Resource Adapter"):
 print "Updating archivePath and classpath of " + ra
 AdminConfig.modify(ra, [['archivePath', "${WAS_INSTALL_ROOT}/installedConnectors/wmq.jmsra.rar"]])
 AdminConfig.unsetAttributes(ra, ['classpath'])
 AdminConfig.modify(ra, [['classpath', "${WAS_INSTALL_ROOT}/installedConnectors/wmq.jmsra.rar"]])
 AdminConfig.save()
 #end if
#end for
```

**Tip:** When saving the file, make sure that it is saved as a python file rather than a text file.

4. Use the WebSphere Application Server wsadmin tool to run the Jython script that you have just created. Open a command prompt and navigate to the \bin directory in the home directory for the WebSphere Application Server, for example C:\Program Files\IBM\WebSphere\AppServer1\bin directory, then type the following command and press Return:

```
wsadmin -lang jython -f convertWMQRA.py
```

If prompted to do so, enter your username and password.

**Note:** If you run the script against a profile that is part of a network deployment configuration, the script updates all profiles that need updating in that configuration. A full resynchronization might be necessary if you have pre-existing configuration file inconsistencies.

5. If you are running in a network deployment configuration, ensure that the node agents are fully re-synchronized. For more information, see [Synchronizing nodes using the wsadmin scripting tool or Adding, managing, and removing nodes](#).
6. Stop all servers in the profile. If the profile is part of a network deployment configuration, also stop any cluster members in the configuration, stop all node agents in the configuration, and stop the deployment manager. If the profile contains an administrative agent, stop the administrative agent.
7. Run the **osgiCfgInit** command from the *profile\_root/bin* directory. The **osgiCfgInit** command resets the class cache used by the OSGi runtime environment. If the profile is part of a network deployment configuration, run the **osgiCfgInit** command from the *profile\_root/bin* directory of every profile that is part of the configuration.
8. Restart all servers in the profile. If the profile is part of a network deployment configuration, also restart any cluster members in the configuration, restart all node agents in the configuration, and restart the deployment manager. If the profile contains an administrative agent, restart the administrative agent.
9. Repeat step 2 to check that the resource adapter is now at the correct level.

## What to do next

If you continue to experience problems after performing the steps described in this topic, and you have previously used the **Update resource adapter** button on the JMS Provider Settings panel in the WebSphere Application Server administrative console to update the IBM MQ resource adapter on any nodes in your environment, it is possible that you are experiencing the issue described in APAR PM10308.

### Related information:

Using the IBM MQ resource adapter

#### **Related information for WebSphere Application Server Version 8.5.5**

Ensuring that servers use the latest available IBM MQ resource adapter maintenance level

[Synchronizing nodes using the wsadmin scripting tool](#)

[Adding, managing, and removing nodes](#)

[JMS provider settings](#)

#### **Related information for WebSphere Application Server Version 8.0**

Ensuring that servers use the latest available IBM MQ resource adapter maintenance level

[Synchronizing nodes using the wsadmin scripting tool](#)

[Adding, managing, and removing nodes](#)

[JMS provider settings](#)

#### **Related information for WebSphere Application Server Version 7.0**

Ensuring that servers use the latest available IBM MQ resource adapter maintenance level

[Synchronizing nodes using the wsadmin scripting tool](#)

[Adding, managing, and removing nodes](#)

[JMS provider settings](#)

## Configuring the JMS PROVIDERVERSION property

The IBM MQ messaging provider has three modes of operation: normal mode, normal mode with restrictions, and migration mode. You can set the JMS **PROVIDERVERSION** property to select which of these modes a JMS application uses to publish and subscribe.



## About this task

The selection of the IBM MQ messaging provider mode of operation can be primarily controlled by setting the **PROVIDERVERSION** connection factory property. The mode of operation can also be selected automatically if a mode has not been specified.

The **PROVIDERVERSION** property differentiates between the three IBM MQ messaging provider modes of operation:

### IBM MQ messaging provider normal mode

Normal mode uses all the features of an IBM MQ queue manager to implement JMS. This mode is optimized to use the JMS 2.0 API and functionality.

### IBM MQ messaging provider normal mode with restrictions

Normal mode with restrictions uses the JMS 2.0 API, but not the new features, that is, shared subscriptions, delayed delivery, and asynchronous send.

### IBM MQ messaging provider migration mode

With migration mode, you can connect to a IBM MQ Version 8.0 queue manager, but none of the features of a IBM WebSphere MQ Version 7.0 or later queue manager, such as read ahead and streaming, are used.




Figure 164. Messaging provider modes

## Procedure

To configure the **PROVIDERVERSION** property for a specific connection factory:

- To configure the **PROVIDERVERSION** property using MQ Explorer, see [Configuring queue managers and objects](#).
- To configure the **PROVIDERVERSION** property using the JMS administration tool, see [Configuring queue managers and objects](#).
- To configure the **PROVIDERVERSION** property in a JMS application using the IBM JMS extensions or IBM MQ JMS extensions, see [Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application](#).

To override connection factory provider mode settings for all connection factories in the JVM:

- To override connection factory provider mode settings, use the `com.ibm.msg.client.wmq.overrideProviderVersion` property. If you cannot change the connection factory that you are using, you can use the `com.ibm.msg.client.wmq.overrideProviderVersion` property to override any setting on the connection factory. This override applies to all connection factories in the JVM but the actual connection factory objects are not modified.

**Related information:**

PROVIDERVERSION

JMS provider version troubleshooting

Connection factory properties

Dependencies between properties of IBM MQ classes for JMS objects

**IBM MQ messaging provider modes of operation**

You can select which of these modes of operation a JMS application uses to publish and subscribe by setting the PROVIDERVERSION property for the connection factory to the appropriate value. In some cases, the PROVIDERVERSION property is set as unspecified, in which case the JMS client uses an algorithm to determine which mode of operation to use.

**PROVIDERVERSION property values**

You can set the connection factory **PROVIDERVERSION** property to any of the following values:

**8 - normal mode**

The JMSApplication uses normal mode. This mode uses all the features of an IBM MQ queue manager to implement JMS.

**7 - normal mode with restrictions**

The JMSApplication uses normal mode with restrictions. This mode uses the JMS 2.0 API, but not the new features such as shared subscriptions, delayed delivery, or asynchronous send.

**6- migration mode**

The JMS application uses migration mode. In migration mode, the IBM MQ classes for JMS use the features and algorithms similar to those that are supplied with IBM WebSphere MQ Version 6.0.

**unspecified (the default value)**

The JMS client uses an algorithm to determine which mode of operation is used.

The value that you specify for the **PROVIDERVERSION** property must be a string. If you are specifying an option of 8, 7 or 6, you can do this in any of the following formats:

- V.R.M.F
- V.R.M
- V.R
- V

where V, R, M and F are integer values greater than or equal to zero. The extra R, M and F values are optional and are available for you to use in case fine grained control is needed. For example, if you wanted to use a **PROVIDERVERSION** level of 7, you could set **PROVIDERVERSION** = 7, 7.0, 7.0.0 or 7.0.0.0.

**Types of connection factory object**

You can set the **PROVIDERVERSION** property for the following types of connection factory object:

- MQConnectionFactory
- MQQueueConnectionFactory
- MQTopicConnectionFactory
- MQXAConnectionFactory
- MQXAQueueConnectionFactory
- MQXAQueueConnectionFactory
- MQXAQueueConnectionFactory
- MQXATopicConnectionFactory

For more information about these different types of connection factory, see “Configuring JMS objects using the administration tool” on page 1159.

**Related information:**

Architecture and overview of features

**PROVIDERVERSION normal mode:**

Normal mode uses all the features of an IBM MQ queue manager to implement JMS. This mode is optimized to use the JMS 2.0 API and functionality.

The following flowchart shows the checks that the JMS client makes to determine whether a normal mode connection can be created.

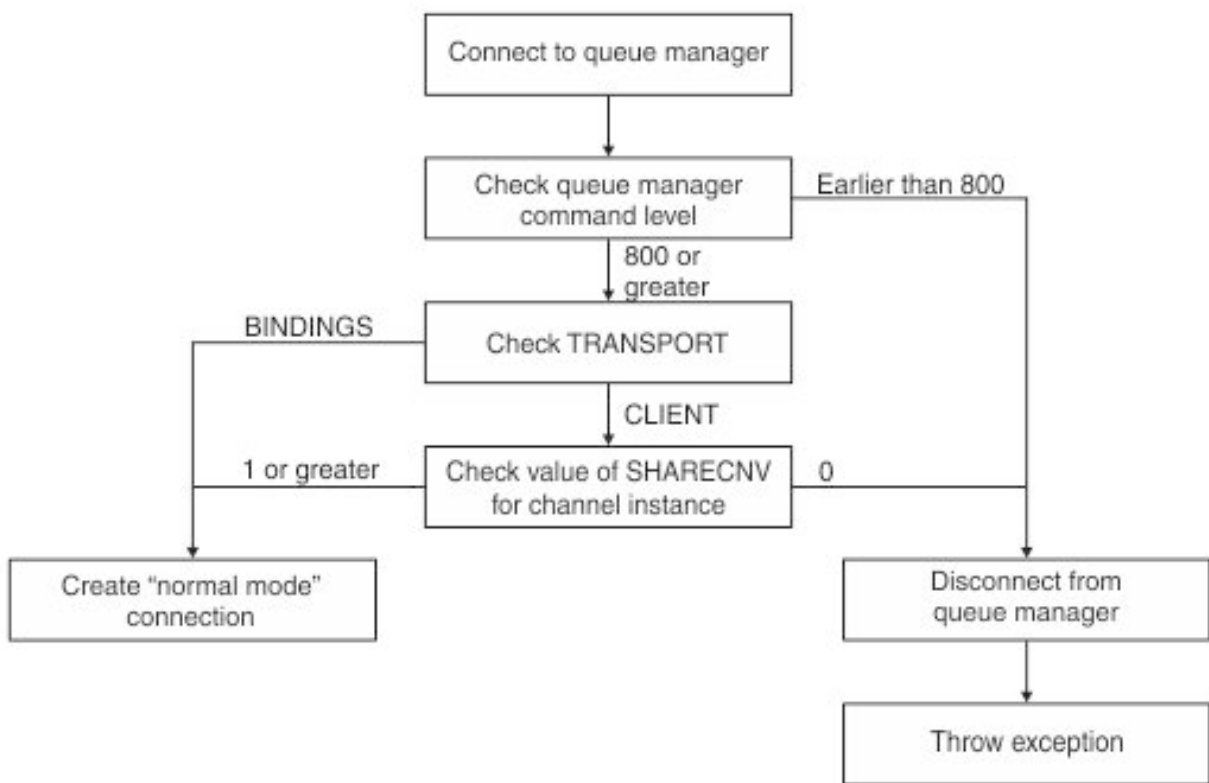


Figure 165. PROVIDERVERSION normal mode

If the queue manager specified in the connection factory settings has a command level of 800 or greater, and the **TRANSPORT** property of the connection factory is set to BINDINGS, a normal mode connection is created without checking any further properties.

If the queue manager specified in the connection factory settings has a command level of 800 or greater, and the **TRANSPORT** property is set to CLIENT, the **SHARECNV** property on the server connection channel is also checked. This check is needed because IBM MQ messaging provider normal mode uses the sharing conversations feature. Therefore, for a normal mode connection attempt to be successful, the **SHARECNV** property, which controls the number of conversations that can be shared, must have a value of 1 or greater.

If all the checks shown in the flowchart are successful, a normal mode connection to the queue manager is created and all of the JMS 2.0 API and features, that is, asynchronous send, delayed delivery, and shared subscription, can then be used.

An attempt to create a normal mode connection fails for either of the following reasons:

- The queue manager specified in the connection factory settings has a command level that is earlier than 800. In this case, the createConnection method fails with an exception JMSFMQ0003.
- The **SHARECNV** property on the server connection channel is set to 0. If this property does not have a value of 1 or greater, the createConnection method fails with an exception JMSSC5007.

**Related information:**

Dependencies between properties of IBM MQ classes for JMS objects

DEFINE CHANNEL (SHARECNV property)

TRANSPORT

**PROVIDERVERSION normal mode with restrictions:**

Normal mode with restrictions uses the JMS 2.0 API, but not the new IBM MQ Version 8.0 features such as shared subscriptions, delayed delivery, or asynchronous send.

The following flowchart shows the checks that the JMS client makes to determine whether a normal mode with restrictions connection can be created .

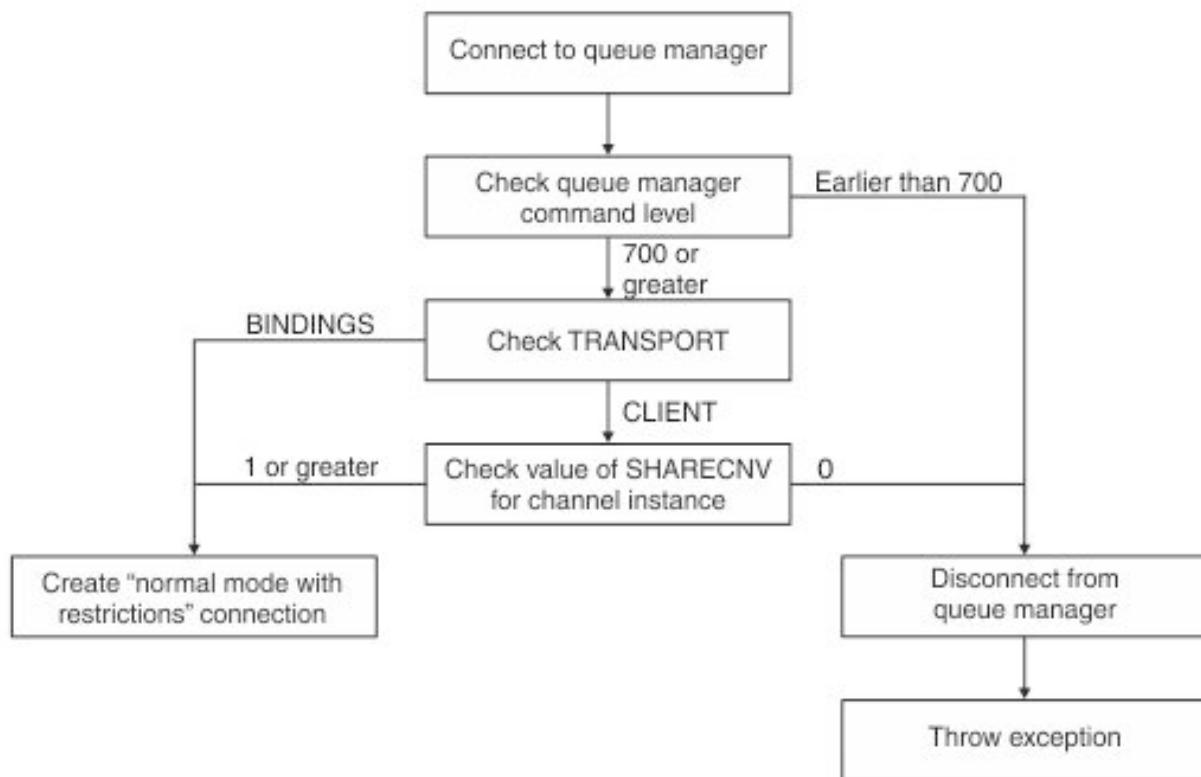


Figure 166. PROVIDERVERSION normal mode with restrictions

If the queue manager specified in the connection factory settings has a command level of 700 or greater, and the **TRANSPORT** property of the connection factory is set to BINDINGS, a normal mode connection is created without checking any further properties.

If the queue manager specified in the connection factory settings has a command level of 700 or greater, and the **TRANSPORT** property is set to CLIENT, the **SHARECNV** property on the server connection channel is also checked. This check is needed because IBM MQ messaging provider normal mode with restrictions uses the sharing conversations feature. Therefore, for a normal mode with restrictions connection attempt to be successful, the **SHARECNV** property, which controls the number of conversations that can be shared, must have a value of 1 or greater.

If all the checks shown in the flowchart are successful, a normal mode with restrictions connection to the queue manager is created and you can then use the JMS 2.0 API, but not the asynchronous send, delayed delivery, or shared subscription features.

An attempt to create a normal mode with restrictions connection fails for either of the following reasons:

- The queue manager specified in the connection factory settings has a command level that is earlier than 700. In this case, the createConnection method fails with exception JMSFCC5008.
- The **SHARECNV** property on the server connection channel is set to 0. If this property does not have a value of 1 or greater, the createConnection method fails with an exception JMSCC5007.

**Related information:**

Dependencies between properties of IBM MQ classes for JMS objects

DEFINE CHANNEL (SHARECNV property)

TRANSPORT

**PROVIDERVERSION migration mode:**

For migration mode, the IBM MQ classes for JMS use the features and algorithms similar to those that are supplied with IBM WebSphere MQ Version 6.0, such as queued publish/subscribe, selection implemented on the client side, non-multiplex channels, and polling used to implement listeners.




Figure 167. PROVIDERVERSION migration mode

If you want to connect to WebSphere Message Broker Version 6.0 or 6.1 using IBM MQ Enterprise Transport Version 6.0, you must use migration mode.

You can connect to an IBM MQ Version 8.0 queue manager using migration mode, but none of the new features of an IBM MQ classes for JMS queue manager are used, for example, read ahead or streaming. If you have an IBM MQ Version 8.0 client connecting to an IBM MQ Version 8.0 queue manager on a distributed platform or an IBM MQ Version 8.0 queue manager on z/OS, then the message selection is done by the queue manager rather than on the client system.

If IBM MQ messaging provider migration mode is specified and the IBM MQ classes for JMS attempt to use any of the JMS 2.0 API, the API method call fails with the exception JMSSC5007.

**Related information:**

Dependencies between properties of IBM MQ classes for JMS objects  
TRANSPORT

**PROVIDERVERSION unspecified:**

When the **PROVIDERVERSION** property of a connection factory is unspecified, the JMS client uses an algorithm to determine which mode of operation is used for connecting to the queue manager. A connection factory that was created in the JNDI namespace with a previous version of IBM MQ classes for JMS takes the unspecified value when the connection factory is used with the new version of IBM MQ classes for JMS.

If the **PROVIDERVERSION** property is unspecified, the algorithm is used when the createConnection method is called. The algorithm checks a number of connection factory properties to determine if IBM MQ messaging provider normal mode, normal mode with restrictions, or IBM MQ messaging provider migration mode is required. Normal mode is always attempted first, and then normal mode with restrictions. If neither of these types of connection can be made, the JMS client disconnects from the queue manager and then connects with the queue manager again to attempt a migration mode connection.

**Checking of BROKERVER, BROKERQGR, PSMODE, and BROKERCONQ properties**

The checking of property values begins with the **BROKERVER** property as shown in Figure 1.

If the **BROKERVER** property is set to V1, the **TRANSPORT** property is checked next as shown in Figure 2. However, if the **BROKERVER** property is set to V2, the additional checking shown in Figure 1 is done before the **TRANSPORT** property is checked.

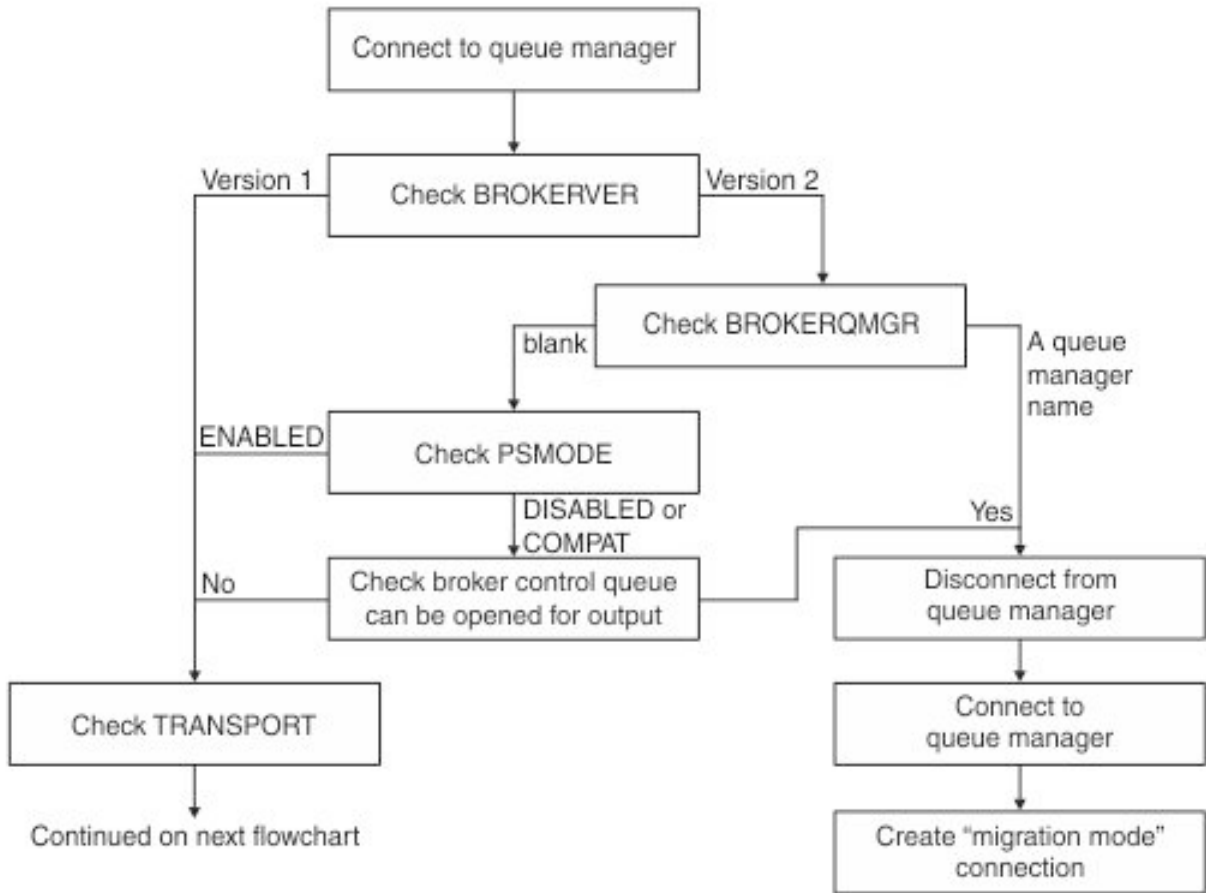


Figure 168. PROVIDERVERSION unspecified

If the **BROKERVER** property is set to V2, for a normal mode connection to be possible, the **BROKERQMGR** property must be blank. Additionally, either the **PSMODE** attribute on the queue manager must be set to ENABLED or the broker control queue specified by the **BROKERCONQ** property must not be able to be opened for output.

If the property values are set as required for a normal mode connection, checking next moves on to the **TRANSPORT** property as shown in Figure 2.

If the property values are not set as required for a normal mode connection, the JMS client disconnects from the queue manager and then reconnects and creates a migration mode connection. This happens in the following cases:

- If the **BROKERQMGR** property is blank and the **PSMODE** attribute on the queue manager is set to COMPAT or DISABLED and the broker control queue specified by the **BROKERCONQ** property can be opened for output (that is, MQOPEN for output succeeds).
- If the **BROKERQMGR** property specifies a queue name.

### Checking of TRANSPORT property and command level

Figure 2 shows the checks that are made for the **TRANSPORT** property and command level of the queue manager.

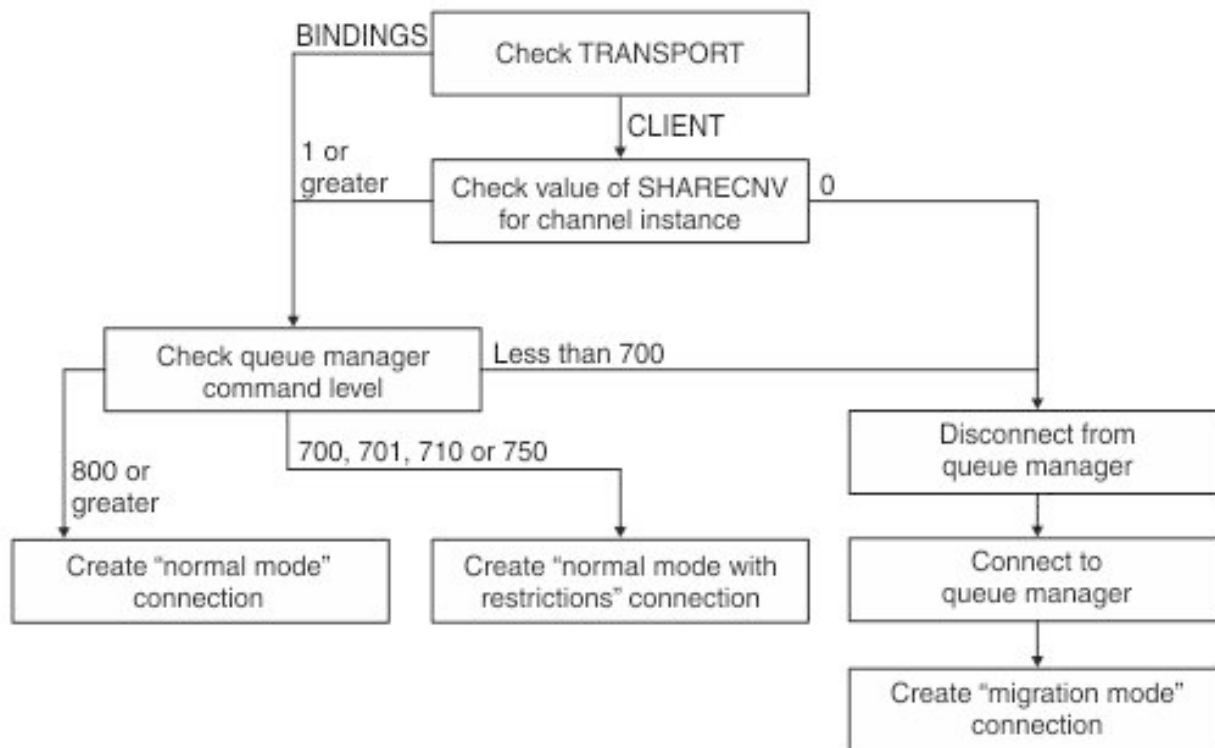


Figure 169. PROVIDERVERSION unspecified (continued)

A normal mode connection is created in either of the following cases:

- The **TRANSPORT** property of the connection factory is set to **BINDINGS**, and the queue manager has a command level of 800 or greater.
- The **TRANSPORT** property is set to **CLIENT**, the **SHARECNV** property on the server connection channel has a value of 1 or greater, and the queue manager has a command level of 800 or greater.

If the queue manager has a command level of 700, 701, 710 or 750, a normal mode with restrictions connection to the queue manager is created.

If the queue manager has a command level of less than 700, the JMS client disconnects from the queue manager and then connects with the queue manager again to create a migration mode connection.

A migration mode connection is also created if the **TRANSPORT** property is set to **CLIENT** and the **SHARECNV** property on the server connection channel has a value of 0.



**Related information:**

Dependencies between properties of IBM MQ classes for JMS objects

ALTER QMGR (PSMODE attribute)

BROKERCONQ

BROKERQMGR

BROKERVER

DEFINE CHANNEL (SHARECNV property)

TRANSPORT

**When to override the PROVIDERVERSION default setting**

If a connection factory that was created in the JNDI namespace with a previous version of IBM MQ classes for JMS is used with the new version of IBM MQ classes for JMS, the **PROVIDERVERSION** property for the connection factory is set to the default value of unspecified and an algorithm is used to determine which IBM MQ messaging provider mode of operation is used. However, there are two cases where you must override the default selection for the **PROVIDERVERSION** property so that the IBM MQ classes for JMS can work correctly.

**Note:** The migration mode that is described in this topic is for migration from IBM WebSphere MQ Version 6.0 to IBM WebSphere MQ Version 7.0. It does not apply to migration from later releases.

IBM WebSphere MQ Version 6.0, WebSphere Application Server Version 6.0.x, and WebSphere Message Broker Version 6 are out of support, and therefore this topic is included only for reference purposes.

When the **PROVIDERVERSION** property is set to the default of unspecified, an algorithm is used to determine which mode of operation to use, as described in “**PROVIDERVERSION** unspecified” on page 1182. However, you cannot use this algorithm in the following two scenarios.

1. If WebSphere Message Broker and WebSphere Event Broker are in compatibility mode, you must specify a value for the **PROVIDERVERSION** property for WebSphere Message Broker and WebSphere Event Broker to work correctly.
2. If you are using WebSphere Application Server Version 6.0.1, WebSphere Application Server Version 6.0.2, or WebSphere Application Server Version 6.1, connection factories are defined by using the WebSphere Application Server administrative console.

In WebSphere Application Server, the default value of the **BROKERVER** property on a connection factory is V2. The default value for the **BROKERVER** property for connection factories that are created by using the JMS administration tool **JMSAdmin** or MQ Explorer is V1. This property is now unspecified in IBM MQ.

If the **BROKERVER** property is set to V2, either because it was created by WebSphere Application Server or the connection factory has been used for publish/subscribe before, and has an existing queue manager that has a **BROKERCONQ** property defined (because it has been used for publish/subscribe messaging before), the IBM MQ messaging provider migration mode is used.

However, if you want the application to use peer-to-peer communication and the application is using an existing queue manager that has ever been used for publish/subscribe, and has a connection factory with **BROKERVER** set to 2, which is the default setting if the connection factory was created in WebSphere Application Server, the IBM MQ messaging provider migration mode is used. Using IBM MQ messaging provider migration mode in this case is unnecessary; use IBM MQ messaging provider normal mode instead. You can use one of the following methods to work around this:

- Set **BROKERVER** to 1 or unspecified. The option that you choose depends on your application.
- Set **PROVIDERVERSION** to 8, or 7, which are custom properties in WebSphere Application Server Version 6.1.

Alternatively, use the client configuration property, or modify the queue manager connected so that it does not have the **BROKERCONQ** property set, or make the queue unusable.

## Configuring provider version information in WebSphere Application Server

To configure provider version information in WebSphere Application Server, you can either use the administrative console or wsadmin commands.

### Procedure

To configure provider version information for an IBM MQ connection factory or activation specification object in WebSphere Application Server, see the *Related information* for links to further information in the WebSphere Application Server product documentation.

#### Related information for WebSphere Application Server Version 8.5.5

IBM MQ messaging provider connection factory settings

**createWMQConnectionFactory** command

IBM MQ messaging provider activation specification settings

**createWMQActivationSpec** command

#### Related information for WebSphere Application Server Version 8.0.0

IBM MQ messaging provider connection factory settings

**createWMQConnectionFactory** command

IBM MQ activation specification settings

**createWMQActivationSpec** command

#### Related information for WebSphere Application Server Version 7.0.0

IBM MQ messaging provider connection factory settings

**createWMQConnectionFactory** command

IBM MQ activation specification settings

**createWMQActivationSpec** command

---

## Configuring HP Integrity NonStop Server

Use this information to help you to configure your IBM MQ HP Integrity NonStop Server client for HP Integrity NonStop Server installation.

For details about configuring a client by using a configuration file, see “Configuring a client using a configuration file” on page 727.

For details about configuring a client by using environment variables, see “Using IBM MQ environment variables” on page 747.

If you are performing IBM MQ HP Integrity NonStop Server client for HP Integrity NonStop Server operations under TMF/Gateway, see the subtopics for information about how to configure the TMF/Gateway. Included are an overview of the Gateway process, configuring the Gateway to run under Pathway, and configuring the client initialization file to enable your IBM MQ HP Integrity NonStop Server client for HP Integrity NonStop Server to reach the TMF Gateway.

This section also contains IBM MQ HP Integrity NonStop Server client for HP Integrity NonStop Server specific information about granting permissions to channels.

## Gateway process overview

The HP NonStop Transaction Management Facility (TMF) provides services to enable a gateway process to register as a resource manager. The IBM MQ for HP Integrity NonStop Server provided TMF/Gateway process runs under Pathway.

IBM MQ for HP Integrity NonStop Server registers a single gateway process for each queue manager that is coordinated by TMF, therefore you must configure a separate TMF/Gateway for each queue manager that is to participate in TMF coordinated units of work. This registration is so that each queue manager is an independent resource manager, and for administrative purposes, registering each queue manager once with HP NonStop TMF results in an easy to understand mapping.

For multiple installations of IBM MQ for HP Integrity NonStop Server, you must nominate a single gateway process from one of these installations for each queue manager to be coordinated by TMF.

The interface to the gateway process supports any client at the same version or earlier.

For more information about administering the gateway process, see *Administering HP Integrity NonStop Server* .

## Configuring Gateway to run under Pathway

TMF/Gateway is the interface between the HP NonStop Transaction Management Facility (TMF) and IBM MQ that enables TMF to be the transaction coordinator for IBM MQ transactions.

The IBM MQ provided TMF/Gateway converts transactions from TMF coordination into eXtended Architecture (XA) transaction coordination to communicate with the remote queue manager.

You must have one TMF/Gateway per queue manager that requires coordination, and client configuration is required so that the client can connect to the correct Gateway.

The TMF/Gateway can use all the mechanisms available to the client to communicate with a queue manager. Configure the TMF/Gateway in the way you would for your other applications.

The TMF/Gateway is not a HP Integrity NonStop Server process pair and is designed to run in a Pathway environment. The TMF/Gateway creates permanent resources within TMF, which it reuses on subsequent runs, therefore the TMF/Gateway must always be run under the same user authority.

## Defining the serverclass

TMF/Gateway is hosted as a serverclass within a Pathway environment. To define the serverclass, you must set the following server attributes:

### **PROCESSTYPE = OSS**

Specifies the type of servers in the serverclass. The Gateway process is a multi-threaded OSS program. This attribute is mandatory, and must be set to OSS.

### **MAXSERVERS = 1**

Specifies the maximum number of server processes in this serverclass that can run at the same time. There can be only a single Gateway process for any queue manager. This attribute is mandatory and must be set to 1.

### **NUMSTATIC = 1**

Specifies the maximum number of static servers within this serverclass. The Gateway process must be run as a static server. This attribute is mandatory and must be set to 1.

### **TMF = ON**

Specifies whether servers in this serverclass can lock and update data files that are audited by the

TMF subsystem. The Gateway process participates in the TMF transactions of IBM MQ client applications therefore this attribute must be set to 0N.

**PROGRAM = <OSS installation path>/opt/mqm/bin/runmqtmf**

For IBM MQ client for IBM MQ, this attribute must be runmqtmf. This attribute must be the absolute OSS path name. Case is significant.

**ARGLIST = -m<QMgr name> [, -c<channel name>] [, -p<port>] [, -h<host name>] [, -n<max threads>]**

These attributes provide parameters to the Gateway process, where:

- QMgrName is the name of the queue manager for this Gateway process. If you are using a queue sharing group (or other port distribution technology), this parameter must be targeted to a specific queue manager. This parameter is mandatory.
- channel name is the name of the server channel on the queue manager to be used by the Gateway process. This parameter is optional.
- port is the TCP/IP port for the queue manager. This parameter is optional.
- host name is the host name for the queue manager. This parameter is optional.
- max threads is the maximum number of worker threads that are created by the Gateway process. This parameter can be a value of 10 or greater. The lowest value that is used is 10 even if a value lower than 10 is specified. If no value is provided, the Gateway process creates up to a maximum of 50 threads.

Use the -c, -p, and -h attributes as an alternative method of providing connection information to the Gateway, in addition to that described in "Configuring the TMF/Gateway using environment variables." If you specify one or more, but not all of the -c, -p, and -h attributes, then those attributes that you do not specify default to the following values:

- channel name defaults to SYSTEM.DEF.SVRCONN
- host name defaults to localhost
- port defaults to 1414

If any of the parameters you supply are invalid, the TMF/Gateway issues diagnostic message AMQ5379 to the error log and terminates.

**OWNER = ID**

The user ID under which the Gateway runs and that must be granted connect authority to the queue manager.

**SECURITY = "value"**

Specifies the users, in relation to the Owner attribute, who can access the Gateway from an IBM MQ client application.

**LINKDEPTH** and **MAXLINKS** must be configured with values appropriate for the expected number of IBM MQ client applications that might want to concurrently communicate with the Gateway. If these values are set too low, you might see occurrences of the error message AMQ5399 issued from client applications.

For more information about these server attributes, see the *HP NonStop TS/MP 2.5 System Management Manual*.

## Configuring the TMF/Gateway using environment variables

One of the most commonly used methods to define the TMF/Gateway is to set the MQSERVER environment variable, for example:

```
ENV MQSERVER=<channel name>/<transport>/<host name>(<listener port>)
```

ENV at the beginning of the command is Pathway notation.

## Configuring the client initialization file

If you are using the HP NonStop Transaction Management Facility (TMF), you must have an IBM MQ client initialization file to enable your IBM MQ client for the HP Integrity NonStop Server to reach the TMF Gateway.

An IBM MQ client initialization file for HP Integrity NonStop Server can be held in a number of locations, for more information, see “Location of the client configuration file” on page 728.

For details of the contents of the configuration file, together with an example, see “Configuring a client using a configuration file” on page 727. Use the TMF and TmfGateway stanzas to specify the TMF queue manager and server details, for more information, see “TMF and TmfGateway stanzas” on page 746.

An example of the entries for a IBM MQ client for HP Integrity NonStop Server is:

```
TMF:
PathMon=$PSD1P
```

```
TmfGateway:
QManager=MQ5B
Server=MQ-MQ5B
```

```
TmfGateway:
QManager=MQ5C
Server=MQ-MQ5C
```

For more information about configuring a client using environment variables, see “Using IBM MQ environment variables” on page 747.

## Granting permissions to channels

Granting permissions to channels on IBM MQ client for HP Integrity NonStop Server is identical to other operating systems, however you must know the identification of the owner that the gateway is running under.

You can then use the identification of the owner of the gateway to grant appropriate permissions. The important difference is that granting permissions to queue manager channels is not under the authority of any application.

Use the `setmqaut` command to both to grant an authorization, that is, give an IBM MQ principal or user group permission to perform an operation, and to revoke an authorization, that is, remove the permission to perform an operation.

---

## Configuring IBM MQ using Docker

► V8.0.0.4

Use this information to configure IBM MQ using Docker.

Docker allows you to package an IBM MQ queue manager or IBM MQ client application, with all of its dependencies, into a standardized unit for software development.

Changes to your application can be deployed to test and staging systems quickly and easily. This feature can be a major benefit to continuous delivery in your enterprise.

## Docker support on Linux systems

> V8.0.0.4

Note the following information if you are using Docker on a Linux system:

- The base image used by the Docker image must use a Linux operating system that is supported.
- The host image, that the Docker container is running on, must use a 3.16 or newer Linux kernel, and must be an operating system supported by IBM MQ.
- You must use the IBM MQ installers to install the product inside the Docker image.
- Only the following packages are supported:
  - MQSeriesRuntime
  - MQSeriesServer
  - MQSeriesClient
  - MQSeriesJava
  - MQSeriesJRE
  - MQSeriesGSKit
  - MQSeriesMsg
  - MQSeriesMan
- The queue manager data directory (`/var/mqm` by default) must be stored on a Docker volume which keeps persistent state.

**Important:** You cannot use the union file system.

You must either mount a host directory as a data volume, or use a data volume container. See [Manage data in containers](#) for more information.

- Running the `mqconfig` command inside the container must pass, to ensure that the container has the correct configuration.
- Applications can only be locally bound to the queue manager (client BINDINGS mode) when running in the same container as that queue manager.
- You must be able to run IBM MQ control commands, such as `endmqm`, within the container.
- You must be able to obtain files and directories from within the container for diagnostic purposes.

## Planning your own IBM MQ queue manager image using Docker

> V8.0.0.4

Use this information to configure IBM MQ using Docker. There are several requirements to consider when running an IBM MQ queue manager in Docker. The sample Docker image provides a way to handle these requirements, but if you want to use your own image, you need to consider how these requirements are handled.

### Process supervision

When you run a Docker container, you are essentially running a single process (PID 1 inside the container), which can later spawn child processes.

If the main process ends, Docker stops the container. An IBM MQ queue manager requires multiple processes to be running in the background.

For this reason, you need to make sure that your main process stays active as long as the queue manager is running. It is good practice to check that the queue manager is active from this process, for example, by performing administrative queries.

## Populating /var/mqm

Docker containers must be configured with /var/mqm as a Docker volume.

When you do this, the directory of the volume is empty when the container first starts. This directory is usually populated at installation time, but installation and runtime are separate environments when using Docker.

To solve this, when your container starts, you can use the `amqicdir` command to populate /var/mqm when it runs for the first time.

## Building a sample IBM MQ queue manager image using Docker

► V8.0.0.4

Use this information to build a sample Docker image for running an IBM MQ queue manager in a Docker container.

Firstly, you will build a base image containing an Ubuntu Linux file system and a clean installation of IBM MQ.

Secondly, you will build another Docker image layer on top of the base, which adds some IBM MQ configuration to allow basic user ID and password security.

Finally, you will run a Docker container using this image as its file system, with the contents of /var/mqm provided by a container-specific Docker volume on the host file system of Docker.

## Building a sample base IBM MQ queue manager image

► V8.0.0.4

In order to use IBM MQ in Docker, you need initially to build a base image with a clean IBM MQ installation. The following steps show you how to build a sample base image, using code hosted on GitHub.

### Procedure

1. Install the prerequisite packages.

These instructions make use of some Linux packages that you must install.

- On Ubuntu:

```
sudo apt-get install python git
```

- On Red Hat Enterprise Linux:

```
sudo yum install python git
```

2. Create a downloads directory by issuing the command `mkdir downloads`.
3. Download the IBM MQ server for Linux image, using Passport Advantage. See Installation using Electronic Software Download for more details.

For example, select the `WS_MQ_V8.0.0.4_LINUX_ON_X86_64_IM.tar.gz` file, and place the file in the downloads directory that you have created.

4. Make the IBM MQ server for Linux image (tar.gz) file available on an HTTP or FTP server.

The reason for this is to save space in the Docker image layers. Every instruction in a Docker file causes a new image layer to be created.

If you use the **ADD** or **COPY** instructions, followed by a **RUN** instruction to install, then the files added or copied will be committed to a new image layer.

Even if you delete the file in subsequent layers, the file still exists in the previous layer. For this reason, it is good practice to download and install within a single **RUN** command, which means the files need to be available on the network.

For example, you can use Python to run an HTTP server, serving all files in your current directory:

```
pushd downloads
nohup python -m SimpleHTTPServer 8000 &
popd
```

5. Extract the sample files, for building a supported Docker image, from GitHub by issuing the following command:

```
git clone -b mq-8 https://github.com/ibm-messaging/mq-docker mq-docker
```

6. Identify your local IP address.

Your address is specific to your local environment, but should be available if you run the following command:

```
ip addr show
```

Note that localhost does not work.

7. Edit the Docker file and change the existing entry for IBM MQ\_URL to point to the file on your local file server. For example:

```
MQ_URL=http://10.0.2.15:8000/WS_MQ_V8.0.0.4_LINUX_ON_X86_64_IM.tar.gz
```

8. Change your directory by issuing the following command:

```
cd mq-docker
```

9. Build the base IBM MQ image by issuing the following command:

```
sudo docker build --tag mq .
```

## Results

You now have a base Docker image with IBM MQ installed.

## Building a sample configured IBM MQ queue manager image

**V8.0.0.4**

Once you have built your generic base IBM MQ Docker image, you need to apply your own configuration to allow secure access. To do this, create your own Docker image, using the generic image as a parent. The following steps show you how to build a sample image, with a minimal security configuration.

## Procedure

1. Create a new directory, and add a file called config.mqsc, with the following contents:

```
DEFINE CHANNEL(PASSWORD.SVRCONN) CHLTYPE(SVRCONN)
SET CHLAUTH(PASSWORD.SVRCONN) TYPE(BLOCKUSER) USERLIST('nobody') +
DESCR('Allow privileged users on this channel')
SET CHLAUTH('*') TYPE(ADDRESSMAP) ADDRESS('*') USERSRC(NOACCESS) DESCR('BackStop rule')
SET CHLAUTH(PASSWORD.SVRCONN) TYPE(ADDRESSMAP) ADDRESS('*') USERSRC(CHANNEL) CHCKCLNT(REQUIRED)
ALTER AUTHINFO(SYSTEM.DEFAULT.AUTHINFO.IDPWOS) AUTHTYPE(IDPWOS) ADOPTCTX(YES)
REFRESH SECURITY TYPE(CONNAUTH)
```

Note that the preceding example uses simple user ID and password authentication. However, you can apply any security configuration that your enterprise requires.

2. Create a file called Dockerfile, with the following contents:

```
FROM mq
RUN useradd johndoe -G mqm && \
 echo johndoe:passwd | chpasswd
COPY config.mqsc /etc/mqm/
```

where:

- johndoe is the user ID that you want to add
- passwd is the original password



3. Build your custom Docker image using the following command:

```
sudo docker build -t mymq .
```

where "." is the directory containing the two files you have just created.

Docker then creates a temporary container using that image, and runs the remaining commands.

The **RUN** command adds a user named johndoe with password passw0rd and the **COPY** command adds the config.mqsc file into a specific location known by the parent image.

4. Run your new customized image to create a new container, with the disk image you have just created.

Your new image layer did not specify any particular command to run, so that has been inherited from the parent image. The entry point of the parent (the code is available on GitHub):

- Creates a queue manager
- Starts the queue manager
- Creates a default listener
- Then runs any MQSC commands from /etc/mqm/config.mqsc.

Issue the following commands to run your new customized image:

```
sudo docker run \
--env LICENSE=accept \
--env MQ_QMGR_NAME=QM1 \
--volume /var/example:/var/mqm \
--publish 1414:1414 \
--detach \
mymq
```

where the:

#### **First env parameter**

Passes an environment variable into the container, which acknowledges your acceptance of the license for IBM WebSphere MQ. You can also set the LICENSE variable to view to view the license.

See IBM MQ license information for further details on IBM MQ licenses.

#### **Second env parameter**

Sets the queue manager name that you are using.

#### **Volume parameter**

Tells the container that whatever MQ writes to /var/mqm should actually be written to /var/example on the host.

This option means that you can easily delete the container later, and still keep any persistent data. This option also makes it easier to view log files.

#### **Publish parameter**

Maps ports on the host system to ports in the container. The container runs by default with its own internal IP address, which means that you need to specifically map any ports that you want to expose.

In this example, that means mapping port 1414 on the host to port 1414 in the container.

#### **Detach parameter**

Runs the container in the background.

## **Results**

You have built a configured docker image and can view running containers using the docker **ps** command. You can view the IBM MQ processes running in your container using the docker **top** command.

**Attention:** If your container is not shown when you use the docker **ps** command the container might have failed. You can see failed containers using the command `docker ps -a`.

The container ID will be shown by using the docker **ps -a** command, and was also printed when you issued the docker **run** command.

You can view the logs of a container using the docker **logs \${CONTAINER\_ID}** command.

A common problem is that **mqconfig** indicates that certain kernel settings on the Docker host are not correct. Kernel settings are shared between the Docker host and containers, and need to be set correctly (see Hardware and software requirements on UNIX and Linux systems).

For example, the maximum number of open files can be set using the command **sysctl fs.file-max=524288**.

---

## Configuring queue managers on z/OS

▶ z/OS

Use these instructions to configure queue managers on IBM MQ for z/OS.

Before you configure IBM MQ, read about the IBM MQ for z/OS concepts in IBM MQ for z/OS concepts.

▶ z/OS Read about how to plan your IBM MQ for z/OS environment in Planning your IBM MQ environment on z/OS .

### Related concepts:

“Creating and managing queue managers on distributed platforms” on page 687

Before you can use messages and queues, you must create and start at least one queue manager and its associated objects.

“Configuring” on page 687

Create one or more queue managers on one or more computers, and configure them on your development, test, and production systems to process messages that contain your business data.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Configuring connections between the server and client” on page 696

To configure the communication links between IBM MQ MQI clients and servers, decide on your communication protocol, define the connections at both ends of the link, start a listener, and define channels.

### Related information:

IBM MQ technical overview

Security

Administering IBM MQ

▶ z/OS Administering IBM MQ for z/OS

Planning

▶ z/OS Issuing commands

▶ z/OS The IBM MQ for z/OS utilities

## Preparing to customize your IBM MQ for z/OS queue managers

Use this topic when customizing your queue managers with details of installable features, national language features, and information about testing, and setting up security.

## Preparing for customization

The Program Directory for WebSphere MQ for z/OS lists the contents of the IBM MQ installation tape, the program and service level information for IBM MQ, and describes how to install IBM MQ for z/OS using the System Modification Program Extended (SMP/E).

When you have installed IBM MQ, you must carry out a number of tasks before you can make it available to users. See the following sections for a description of these tasks:

- “Customizing IBM MQ for z/OS” on page 1198
- “Testing your queue manager on z/OS” on page 1252
- Setting up security on z/OS

If you are migrating from a previous version of IBM MQ for z/OS, you do not need to perform most of the customization tasks. See Migrating and upgrading IBM MQ for more information about the tasks you must perform.

### Installable features of IBM MQ for z/OS

IBM MQ for z/OS comprises the following features:

- Base** This is required; it comprises all the main functions, including
- Administration and utilities
  - Support for CICS, IMS, and batch type applications using the IBM MQ Application Programming Interface, or C++
  - Distributed queuing facility (supporting both TCP/IP and APPC communications)

#### National language features

These contain error messages and panels in all the supported national languages. Each language has a language letter associated with it. The languages and letters are:

- |          |                           |
|----------|---------------------------|
| <b>C</b> | Simplified Chinese        |
| <b>E</b> | U.S. English (mixed case) |
| <b>F</b> | French                    |
| <b>K</b> | Japanese                  |
| <b>U</b> | U.S. English (uppercase)  |

You must install the US English (mixed case) option. You can also install one or more other languages. (The installation process for other languages requires US English (mixed case) to be installed, even if you are not going to use US English (mixed case).)

#### IBM MQ for z/OS Unix System Services Components

This feature is optional. Select this feature if you want to build and run Java applications that use the Java Message Service (JMS) to connect to IBM MQ for z/OS or if you want to build and run HTTP applications which use HTTP to connect to IBM MQ for z/OS.

### Libraries that exist after installation

IBM MQ is supplied with a number of separate load libraries. Table 144 on page 1196 shows the libraries that might exist after you have installed IBM MQ.

Table 144. IBM MQ libraries that exist after installation

Name	Description
thlqual.SCSQANLC	Contains the load modules for the Simplified Chinese version of IBM MQ.
thlqual.SCSQANLE	Contains the load modules for the U.S. English (mixed case) version of IBM MQ.
thlqual.SCSQANLF	Contains the load modules for the French version of IBM MQ.
thlqual.SCSQANLK	Contains the load modules for the Japanese version of IBM MQ.
thlqual.SCSQANLU	Contains the load modules for the U.S. English (uppercase) version of IBM MQ.
thlqual.SCSQASMS	Contains source for assembler sample programs.
thlqual.SCSQAUTH	The main repository for all IBM MQ product load modules; it also contains the default parameter module, CSQZPARM. This library must be APF-authorized and in PDS-E format.
thlqual.SCSQCICS	Contains extra load modules that must be included in the CICS DFHRPL concatenation. This library must be APF-authorized and in PDS-E format.
thlqual.SCSQCLST	Contains CLISTs used by the sample programs.
thlqual.SCSQCOBC	Contains COBOL copybooks, including copybooks required for the sample programs.
thlqual.SCSQCOBS	Contains source for COBOL sample programs.
thlqual.SCSQCPPS	Contains source for C++ sample programs.
thlqual.SCSQC37S	Contains source for C sample programs.
thlqual.SCSQC370	Contains C headers, including headers required for the sample programs.
thlqual.SCSQDEFS	Contains side definitions for C++ and the Db2 DBRMs for shared queuing.
thlqual.SCSQEXEC	Contains REXX executable files to be included in the SYSEXEC or SYSPROC concatenation if you are using the IBM MQ operations and control panels.
thlqual.SCSQHPPS	Contains header files for C++.
thlqual.SCSQINST	Contains JCL for installation jobs.
thlqual.SCSQLINK	Early code library. Contains the load modules that are loaded at system initial program load (IPL). The library must be APF-authorized.
thlqual.SCSQLOAD	Load library. Contains load modules for non-APF code, user exits, utilities, samples, installation verification programs, and adapter stubs. The library does not need to be APF-authorized and does not need to be in the link list. This library must be in PDS-E format.
thlqual.SCSQMACS	Contains Assembler macros including: sample macros, product macros, and system parameter macros.
thlqual.SCSQMAPS	Contains CICS mapsets used by sample programs.
thlqual.SCSQMSGC	Contains ISPF messages to be included in the ISPMLIB concatenation if you are using the Simplified Chinese language feature for the IBM MQ operations and control panels.
thlqual.SCSQMSGE	Contains ISPF messages to be included in the ISPMLIB concatenation if you are using the U.S. English (mixed case) language feature for the IBM MQ operations and control panels.

Table 144. IBM MQ libraries that exist after installation (continued)

Name	Description
thlqual.SCSQMSGF	Contains ISPF messages to be included in the ISPMLIB concatenation if you are using the French language feature for the IBM MQ operations and control panels.
thlqual.SCSQMSGK	Contains ISPF messages to be included in the ISPMLIB concatenation if you are using the Japanese language feature for the IBM MQ operations and control panels.
thlqual.SCSQMSGU	Contains ISPF messages to be included in the ISPMLIB concatenation if you are using the U.S. English (uppercase) language feature for the IBM MQ operations and control panels.
thlqual.SCSQMVR1	Contains the load modules for distributed queuing. This library must be APF-authorized and in PDS-E format.
thlqual.SCSQPLIC	Contains PL/I include files.
thlqual.SCSQPLIS	Contains source for PL/I sample programs.
thlqual.SCSQPMLA	Contains IPCS panels, for the dump formatter, to be included in the ISPPLIB concatenation. Also contains panels for IBM MQ sample programs.
thlqual.SCSQPMLC	Contains ISPF panels to be included in the ISPPLIB concatenation if you are using the Simplified Chinese language feature for the IBM MQ operations and control panels.
thlqual.SCSQPMLD	Contains ISPF panels to be included in the ISPPLIB concatenation if you are using the U.S. English (mixed case) language feature for the IBM MQ operations and control panels.
thlqual.SCSQPMLF	Contains ISPF panels to be included in the ISPPLIB concatenation if you are using the French language feature for the IBM MQ operations and control panels.
thlqual.SCSQPMLK	Contains ISPF panels to be included in the ISPPLIB concatenation if you are using the Japanese language feature for the IBM MQ operations and control panels.
thlqual.SCSQPMLU	Contains ISPF panels to be included in the ISPPLIB concatenation if you are using the U.S. English (uppercase) language feature for the IBM MQ operations and control panels.
thlqual.SCSQPROC	Contains sample JCL and default system initialization data sets.
thlqual.SCSQSNLC	Contains the load modules for the Simplified Chinese versions of the IBM MQ modules that are required for special purpose function (for example the early code).
thlqual.SCSQSNLE	Contains the load modules for the U.S. English (mixed case) versions of the IBM MQ modules that are required for special purpose function (for example the early code).
thlqual.SCSQSNLF	Contains the load modules for the French versions of the IBM MQ modules that are required for special purpose function (for example the early code).
thlqual.SCSQSNLK	Contains the load modules for the Japanese versions of the IBM MQ modules that are required for special purpose function (for example the early code).
thlqual.SCSQSNLU	Contains the load modules for the U.S. English (uppercase) versions of the IBM MQ modules that are required for special purpose function (for example the early code).

Table 144. IBM MQ libraries that exist after installation (continued)

Name	Description
thlqual.SCSQTBLC	Contains ISPF tables to be included in the ISPTLIB concatenation if you are using the Simplified Chinese language feature for the IBM MQ operations and control panels.
thlqual.SCSQTBLE	Contains ISPF tables to be included in the ISPTLIB concatenation if you are using the U.S. English (mixed case) language feature for the IBM MQ operations and control panels.
thlqual.SCSQTBLF	Contains ISPF tables to be included in the ISPTLIB concatenation if you are using the French language feature for the IBM MQ operations and control panels.
thlqual.SCSQTBLK	Contains ISPF tables to be included in the ISPTLIB concatenation if you are using the Japanese language feature for the IBM MQ operations and control panels.
thlqual.SCSQTBLU	Contains ISPF tables to be included in the ISPTLIB concatenation if you are using the U.S. English (uppercase) language feature for the IBM MQ operations and control panels.

**Note:** Do not modify or customize any of these libraries. If you want to make changes, copy the libraries and make your changes to the copies.

**Related concepts:**

“Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

“Using IBM MQ with IMS” on page 1294

The IBM MQ -IMS adapter, and the IBM MQ - IMS bridge are the two components which allow IBM MQ to interact with IMS.

“Using IBM MQ with CICS” on page 1303

To use IBM MQ with CICS, you must configure the IBM MQ CICS adapter and, optionally, the IBM MQ CICS bridge components.

“Using OTMA exits in IMS” on page 1305

Use this topic if you want to use IMS Open Transaction Manager Access exits with IBM MQ for z/OS.

**Related reference:**

“Upgrading and applying service to Language Environment or z/OS Callable Services” on page 1303

The actions you must take vary according to whether you use CALLLIBS or LINK, and your version of SMP/E.

**Related information:**

IBM MQ for z/OS concepts

Administering IBM MQ for z/OS

## Customizing IBM MQ for z/OS

Use this topic as a step by step guide for customizing your IBM MQ system.

This topic leads you through the various stages of customizing IBM MQ after you have successfully installed it. The installation process is described in the Program Directory, available to download from the IBM Publications Center.

Samples are supplied with IBM MQ to help you with your customization. The sample data set members have names beginning with the four characters CSQ4 and are in the library thlqual.SCSQPROC.

Before you perform the customization tasks described in this topic, there are a number of configuration options that you must consider because they affect the performance and resource requirements of IBM MQ for z/OS. For example, you must decide which globalization libraries you want to use.

## Configuration options

For more information about these options, see *Planning on z/OS*.

The description of each task in this section indicates whether:

- The task is part of the process of customizing IBM MQ. That is, you perform the task once when you customize IBM MQ on the z/OS system. (In a parallel sysplex, you must perform the task for each z/OS system in the sysplex, and ensure that each z/OS system is set up identically.)
- The task is part of adding a queue manager. That is, you perform the task once for each queue manager when you add that queue manager.
- You need to perform the task when migrating. If you are migrating from a previous version of IBM MQ for z/OS, you might not need to perform all these tasks.

Review the tasks when you apply corrective maintenance to IBM MQ and when you install a new version or release of IBM MQ.

None of the tasks require you to perform an IPL of your z/OS system, if you use commands to change the various z/OS system parameters, and perform “Task 12: Update SYS1.PARMLIB members” on page 1215 as suggested.

To simplify operations and to aid with problem determination, ensure that all z/OS systems in a sysplex are set up identically, so that queue managers can be quickly created on any system in an emergency.

For ease of maintenance, consider defining aliases to refer to your IBM MQ libraries; for more information, see *Using an alias to refer to an IBM MQ library*.

## Identify the national language support libraries

You need to specify the appropriate globalization libraries in the JCL that you want to use with IBM MQ (as described in the following sections). Each language is identified by a language letter:

- C** Simplified Chinese
- E** U.S. English (mixed case)
- F** French
- K** Japanese
- U** U.S. English (uppercase)

*Table 145. National language feature libraries*

Description	Japanese	Simplified Chinese	U.S. English (mixed case)	U.S. English (uppercase)	French
Load modules	thlqual.SCSQANLK	thlqual.SCSQANLC	thlqual.SCSQANLE	thlqual.SCSQANLU	thlqual.SCSQANLF
ISPF messages	thlqual.SCSQMSGK	thlqual.SCSQMSGC	thlqual.SCSQMSGE	thlqual.SCSQMSGU	thlqual.SCSQMSGF
ISPF panels	thlqual.SCSQPNLK	thlqual.SCSQPNLC	thlqual.SCSQPNLE	thlqual.SCSQPNLU	thlqual.SCSQPNLF
Special purpose function (for example, early code)	thlqual.SCSQSNLK	thlqual.SCSQSNLC	thlqual.SCSQSNLE	thlqual.SCSQSNLU	thlqual.SCSQSNLF
ISPF tables	thlqual.SCSQTBLK	thlqual.SCSQTBLC	thlqual.SCSQTBLE	thlqual.SCSQTBLU	thlqual.SCSQTBLF

## Customization summary

The following table lists all the steps required to customize IBM MQ for z/OS. It also indicates the following:

- Whether the step has to be performed once only, or repeated for each queue manager.
- Whether you need to repeat the step for each queue-sharing group, or omit the step if you are not using queue-sharing groups.
- Whether the step is required if you are migrating from a previous version of IBM MQ. Some steps might be needed, depending on what you decide about data set and queue manager names; these steps are marked 'Review'.

Table 146. Customization summary

Task	Required when migrating	Repeat for each queue manager	Queue-sharing groups
"Task 1: Identify the z/OS system parameters" on page 1201	Review	-	-
"Task 2: APF authorize the IBM MQ load libraries" on page 1202	Review	-	-
"Task 3: Update the z/OS link list and LPA" on page 1203	Review	-	-
"Task 4: Update the z/OS program properties table" on page 1205	-	-	-
"Task 5: Define the IBM MQ subsystem to z/OS" on page 1205	-	X	-
"Task 6: Create procedures for the IBM MQ queue manager" on page 1209	Review	X	-
"Task 7: Create procedures for the channel initiator" on page 1210	Review	X	-
"Task 8: Define the IBM MQ subsystem to a z/OS WLM service class" on page 1211	-	X	-
"Task 9: Set up the Db2 environment" on page 1211	Review	-	Omit if not used
"Task 10: Set up the coupling facility" on page 1213	Review	-	Repeat for each
"Task 11: Implement your ESM security controls" on page 1214	Review	X	X
"Task 12: Update SYS1.PARMLIB members" on page 1215	Review	-	-
"Task 13: Customize the initialization input data sets" on page 1215	X	X	-
"Task 14: Create the bootstrap and log data sets" on page 1218	-	X	-
"Task 15: Define your page sets" on page 1219	-	X	-
"Task 16: Add the IBM MQ entries to the Db2 tables" on page 1220	Review	X	Repeat for each
"Task 17: Tailor your system parameter module" on page 1221	X	X	-
"Task 18: Tailor the channel initiator parameters" on page 1242	X	X	-



Table 146. Customization summary (continued)

Task	Required when migrating	Repeat for each queue manager	Queue-sharing groups
"Task 19: Set up Batch, TSO, and RRS adapters" on page 1244	Review	-	-
"Task 20: Set up the operations and control panels" on page 1245	Review	-	-
"Task 21: Include the IBM MQ dump formatting member" on page 1246	X	-	-
"Task 22: Suppress information messages" on page 1247	-	-	-
"Task 23: Create procedures for Advanced Message Security" on page 1248	Review	X	-
"Task 24: Set up the started task user Advanced Message Security" on page 1248	Review	X	-
"Task 25: Grant RACDCERT permissions to the security administrator for Advanced Message Security" on page 1251	-	-	-
"Task 26: Grant users resource permissions for IBM MQ Advanced Message Security" on page 1251	-	-	-

**Related concepts:**

"Setting up communications with other queue managers" on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

"Using IBM MQ with IMS" on page 1294

The IBM MQ -IMS adapter, and the IBM MQ - IMS bridge are the two components which allow IBM MQ to interact with IMS.

"Using IBM MQ with CICS" on page 1303

To use IBM MQ with CICS, you must configure the IBM MQ CICS adapter and, optionally, the IBM MQ CICS bridge components.

"Using OTMA exits in IMS" on page 1305

Use this topic if you want to use IMS Open Transaction Manager Access exits with IBM MQ for z/OS.

**Related reference:**


"Upgrading and applying service to Language Environment or z/OS Callable Services" on page 1303

The actions you must take vary according to whether you use CALLLIBS or LINK, and your version of SMP/E.

**Related information:**

IBM MQ for z/OS concepts

Administering IBM MQ for z/OS

 Program Directory for IBM MQ for z/OS

**Task 1: Identify the z/OS system parameters**

Some of the tasks involve updating the z/OS system parameters. You need to know which ones were specified when the system IPL was performed.

- You need to perform this task once for each z/OS system where you want to run IBM MQ.
- You might need to perform this task when migrating from a previous version.

SYS1.PARMLIB(IEASYSpp) contains a list of parameters that point to other members of SYS1.PARMLIB (where pp represents the z/OS system parameter list that was used to perform an IPL of the system).

The entries you need to find are:

**For “Task 2: APF authorize the IBM MQ load libraries”:**

PROG=xx or APF=aa point to the Authorized Program Facility (APF) authorized library list (member PROGxx or IEFAPFaa)

**For “Task 3: Update the z/OS link list and LPA” on page 1203:**

LNK=kk points to the link list (member LNKLSTkk) LPA=mm points to the LPA list (member LPALSTmm)

**For “Task 4: Update the z/OS program properties table” on page 1205:**

SCH=xx points to the Program Properties Table (PPT) (member SCHEDxx)

**For “Task 5: Define the IBM MQ subsystem to z/OS” on page 1205:**

SSN=ss points to the defined subsystem list (member IEFSSNss)

**Task 2: APF authorize the IBM MQ load libraries**

APF-authorize various libraries. Some load modules might already be authorized.

- You need to perform this task once for each z/OS system where you want to run IBM MQ.
- If you are using queue-sharing groups, you must ensure that the settings for IBM MQ are identical on each z/OS system in the sysplex.
- You might need to perform this task when migrating from a previous version.

The IBM MQ load libraries thlqual.SCSQAUTH and thlqual.SCSQLINK must be APF-authorized. You must also APF-authorize the libraries for your national language feature (thlqual.SCSQANLx and thlqual.SCSQSNLx) and for the distributed queuing feature (thlqual.SCSQMVR1). If you are using IBM MQ Advanced Message Security you must also APF authorize the library thlqual.SDRQAUTH.

However, all load modules in the LPA are automatically APF-authorized. So are all members of the link list if the SYS1.PARMLIB member IEASYSpp contains the statement:

```
LNKAUTH=LNKLST
```

LNKAUTH=LNKLST is the default if LNKAUTH is not specified.

Depending on what you choose to put in the LPA or linklist (see “Task 3: Update the z/OS link list and LPA” on page 1203 ), you might not need to put the libraries in the APF link list

**Note:** You must APF-authorize all the libraries that you include in the IBM MQ STEPLIB. If you put a library that is not APF-authorized in the STEPLIB, the whole library concatenation loses its APF authorization.

The APF lists are in the SYS1.PARMLIB member PROGxx or IEAAPFaa. The lists contain the names of APF authorized z/OS libraries. The order of the entries in the lists is not significant. See the *MVS Initialization and Tuning Reference* manual for information about APF lists.

For more information about tuning your system, see SupportPac MP16

If you use PROGxx members with dynamic format, you need only issue the z/OS command SETPROG APF,ADD,DSNAME=h1q.SCSQ XXXX,VOLUME= YYYYYY for the changes to take effect: Where XXXX varies by the library name and where YYYYYY is the volume. Otherwise, if you use static format or IEAAPFaa members, you must perform an IPL on your system.

Note that you must use the actual name of the library in the APF list. If you attempt to use the data set alias of the library, authorization fails.

**Related concepts:**

“Task 3: Update the z/OS link list and LPA”

Update the LPA libraries with the new version of early code libraries. Other code can go in the link list or the LPA.

“Preparing to customize your IBM MQ for z/OS queue managers” on page 1194

Use this topic when customizing your queue managers with details of installable features, national language features, and information about testing, and setting up security.

### **Task 3: Update the z/OS link list and LPA**

Update the LPA libraries with the new version of early code libraries. Other code can go in the link list or the LPA.

- You need to perform this task once for each z/OS system where you want to run IBM MQ.
- If you are using queue-sharing groups, you must ensure that the settings for IBM MQ are identical on each z/OS system in the sysplex.
- You might need to perform this task when migrating from a previous version. For details, see the Program Directory for WebSphere MQ for z/OS.

**Note:** The data set for LPA is version specific. If you are using an existing LPA in the system, contact your system administrator to decide which LPA to use.

### **Early code**

Some IBM MQ load modules need to be added to MVS for IBM MQ to act as a subsystem. These modules are known as the Early code, and they can be executed even if a queue manager is not active. For example, when an operator command is issued on the console with an IBM MQ command prefix, this Early code will get control and check if it needs to start a queue manager, or to pass the request to a running queue manager. This code is loaded into the Link Pack Area (LPA). There is one set of Early modules, which are used for all queue managers, and these need to be at the highest level of IBM MQ. Early code from a higher version of IBM MQ will work with a queue manager with a lower version of IBM MQ, but not the opposite.

#### **IBM MQ**

The early code comprises the following load modules:

- CSQ3INI and CSQ3EPX in the library thqual.SCSQLINK
- CSQ3ECMX in the library thqual.SCSQSNL *x*, where *x* is your language letter.

IBM MQ includes a user modification that moves the contents of the thqual.SCSQSNL *x* library into the thqual.SCSQLINK and informs SMP/E. This user modification is called CSQ8UERL and is described in the Program Directory for WebSphere MQ for z/OS.

When you have updated the early code in the LPA libraries, it is available from the next z/OS IPL (with the CLPA option) to all queue manager subsystems added during IPL from definitions in IEFSSNss members in SYS1.PARMLIB.

You can make it available immediately without an IPL for any new queue manager subsystem added later (as described in “Task 5: Define the IBM MQ subsystem to z/OS” on page 1205 ) by adding it to the LPA as follows:

- If you did not use CSQ8UERL, issue these z/OS commands:  
SETPROG LPA,ADD,MODNAME=(CSQ3INI,CSQ3EPX),DSNAME=thqual.SCSQLINK  
SETPROG LPA,ADD,MODNAME=(CSQ3ECMX),DSNAME=thqual.SCSQSNL *x*

- If you did use CSQ8UERL, you can load the early code into the LPA using the following z/OS command:

```
SETPROG LPA,ADD,MASK=*,DSNAME=thqua1.SCSQLINK
```

- If you are using IBM MQ Advanced Message Security you must also issue the following z/OS command to include an additional module in the LPA:

```
SETPROG LPA,ADD,MODNAME=(CSQ0DRTM),DSNAME=thqua1.SCSQLINK
```

If you have applied maintenance, or you intend to restart a queue manager with a later version or release of IBM MQ, the early code can be made available to queue manager subsystems that are already defined, provided the level of the early code when the z/OS system was started was at least that of Version 5.3.

To make it available, use the following steps:

1. Add it to the LPA using z/OS SETPROG commands as described previously in this topic.
2. Stop the queue manager, using the IBM MQ command STOP QMGR.
3. Ensure that the qmgr.REFRESH.QMGR security profile is set up. See MQSC commands, profiles, and their access levels.
4. Refresh the early code for the queue manager using the IBM MQ command REFRESH QMGR TYPE(EARLY).
5. Restart the queue manager, using the IBM MQ command START QMGR.

The IBM MQ commands STOP QMGR, REFRESH QMGR, and START QMGR are described in MQSC commands.

## Other code

All the IBM MQ supplied load modules in the following libraries are reentrant and can be placed in the LPA:

- SCSQAUTH
- SCSQANL *x*, where *x* is your language letter
- SCSQMVR1

**Important:** However, if you place the libraries in the LPA, whenever you apply maintenance, you have to copy any changed modules manually into the LPA. Therefore, it is preferable to put the IBM MQ load libraries in the link list, which can be updated after maintenance by issuing the z/OS command REFRESH LLA.

This is particularly recommended for SCSQAUTH so that you do not have to include it in several STEPLIBs. Only one language library, SCSQANL *x* should be placed in the LPA or link list. The link list libraries are specified in an LNKLSTkk member of SYS1.PARMLIB.

The distributed queuing facility and CICS bridge (but not the queue manager itself) need access to the Language Environment (LE) runtime library SCEERUN. If you use either of these facilities, you need to include SCEERUN in the link list.

**Related concepts:**

“Task 4: Update the z/OS program properties table”

Some additional PPT entries are needed for the IBM MQ queue manager.

**Task 4: Update the z/OS program properties table**

Some additional PPT entries are needed for the IBM MQ queue manager.

- *You must perform this task once for each z/OS system where you want to run IBM MQ.*
- *If you are using queue-sharing groups, you must ensure that the settings for IBM MQ are identical on each z/OS system in the sysplex.*
- *You do not need to perform this task when migrating from a previous version.*
- *You do need to perform the CSQ0DSRV part of this task when you require IBM MQ Advanced Message Security.*

A sample containing all the required PPT entries is provided in thlqual.SCSQPROC(CSQ4SCHD). Ensure that the required entries are added to the PPT, which you can find in SYS1.PARMLIB(SCHEDxx).

In z/OS 1.12 and later versions, CSQYASCP is already defined to the operating system with the attributes detailed and no longer needs to be included in a SCHEDxx member of PARMLIB.

The IBM MQ queue manager controls swapping itself. However, if you have a heavily-loaded IBM MQ network and response time is critical, it might be advantageous to make the IBM MQ channel initiator nonswappable, by adding the CSQXJST PPT entry, at the risk of affecting the performance of the rest of your z/OS system.

If you require IBM MQ Advanced Message Security, add the CSQ0DSRV PPT entry.

Issue the z/OS command SET SCH= for these changes to take effect.

**Related concepts:**

“Task 5: Define the IBM MQ subsystem to z/OS”

Update the subsystem name table and decide on a convention for command prefix strings.

**Task 5: Define the IBM MQ subsystem to z/OS**

Update the subsystem name table and decide on a convention for command prefix strings.

Repeat this task for each IBM MQ queue manager. You do not need to perform this task when migrating from a previous version.

**Related concepts:**

“Task 6: Create procedures for the IBM MQ queue manager” on page 1209

Each IBM MQ subsystem needs a cataloged procedure to start the queue manager. You can create your own or use the IBM-supplied procedure library.

**Updating the subsystem name table:**

When defining the IBM MQ subsystem you must add an entry to the subsystem name table.

The subsystem name table of z/OS, which is taken initially from the SYS1.PARMLIB member IEFSSNss, contains the definitions of formally defined z/OS subsystems. To define each IBM MQ subsystem, you must add an entry to this table, either by changing the IEFSSNss member of SYS1.PARMLIB, or, preferably, by using the z/OS command SETSSI.

IBM MQ subsystem initialization supports parallel processing, so IBM MQ subsystem definition statements can be added both above and below the BEGINPARALLEL keyword in the IEFSSNss table available at z/OS V1.12 and later.

If you use the SETSSI command, the change takes effect immediately, and there is no need to perform an IPL of your system. Ensure you update SYS1.PARMLIB as well, as described in “Task 12: Update SYS1.PARMLIB members” on page 1215 so that the changes remain in effect after subsequent IPLs.

The SETSSI command to dynamically define an IBM MQ subsystem is:

```
SETSSI ADD,S=ssid,I=CSQ3INI,P='CSQ3EPX,cpf,scope'
```

The corresponding information in IEFSSNss can be specified in one of two ways:

- The keyword parameter form of the IBM MQ subsystem definition in IEFSSNss. This is the recommended method.

```
SUBSYS SUBNAME(ssid) INITRTN(CSQ3INI) INITPARM('CSQ3EPX,cpf,scope')
```

- The positional parameter form of the IBM MQ subsystem definition.

```
ssid,CSQ3INI,'CSQ3EPX,cpf,scope'
```

Do not mix the two forms in one IEFSSNss member. If different forms are required, use a separate IEFSSNss member for each type, adding the SSN operand of the new member to the IEASYSpp SYS1.PARMLIB member. To specify more than one SSN, use SSN=(aa,bb,...) in IEASYSpp.

In the examples,

- ssid** The subsystem identifier. It can be up to four characters long. All characters must be alphanumeric (uppercase A through Z, 0 through 9), it must start with an alphabetic character. The queue manager will have the same name as the subsystem, therefore you can use only characters that are allowed for both z/OS subsystem names and IBM MQ object names.
- cpf** The command prefix string (see “Defining command prefix strings (CPFs)” on page 1207 for information about CPFs).
- scope** The system scope, used if you are running in a z/OS sysplex (see “CPFs in a sysplex environment” on page 1208 for information about system scope).

Figure 170 shows several examples of IEFSSNss statements.

```
CSQ1,CSQ3INI,'CSQ3EPX,+mqs1cpf,S'
CSQ2,CSQ3INI,'CSQ3EPX,+mqs2cpf,S'
CSQ3,CSQ3INI,'CSQ3EPX,++,S'
```

*Figure 170. Sample IEFSSNss statements for defining subsystems*

**Note:** When you have created objects in a subsystem, you cannot change the subsystem name or use the page sets from one subsystem in another subsystem. To do either of these, you must unload all the objects and messages from one subsystem and reload them into another.

Table 147 on page 1207 gives a number of examples showing the associations of subsystem names and command prefix strings (CPFs), as defined by the statements in Figure 170.

Table 147. Subsystem name to CPF associations

IBM MQ subsystem name	CPF
CSQ1	+mqs1cpf
CSQ2	+mqs2cpf
CSQ3	++

**Note:** The ACTIVATE and DEACTIVATE functions of the z/OS command SETSSI are not supported by IBM MQ.

To check the status of the changes, issue the following command in SDSF: /D SSI,L. You will see the new subsystems created with ACTIVE status.

### Defining command prefix strings (CPFs):

Each subsystem instance of IBM MQ can have a command prefix string to identify that subsystem.

Adopt a system-wide convention for your CPFs for all subsystems to avoid conflicts. Adhere to the following guidelines:

- Define a CPF as string of up to eight characters.
- Do not use a CPF that is already in use by any other subsystem, and avoid using the JES backspace character defined on your system as the first character of your string.
- Define your CPF using characters from the set of valid characters listed in Table 149 on page 1208.
- Do not use a CPF that is an abbreviation for an already defined process or that might be confused with command syntax. For example, a CPF such as 'D' conflicts with z/OS commands such as DISPLAY. To avoid this happening, use one of the special characters (shown in Table 149 on page 1208 ) as the first or only character in your CPF string.
- Do not define a CPF that is either a subset or a superset of an existing CPF. For an example, see Table 148.

Table 148. Example of CPF subset and superset rules

Subsystem name	CPF defined	Commands routed to
MQA	!A	MQA
MQB	!B	MQB
MQC1	!C1	MQC1
MQC2	!C2	MQC2
MQB1	!B1	MQB

Commands intended for subsystem MQB1 (using CPF !B1) are routed to subsystem MQB because the CPF for this subsystem is !B, a subset of !B1. For example, if you entered the command:

```
!B1 START QMGR
```

subsystem MQB receives the command:

```
1 START QMGR
```

(which, in this case, it cannot deal with).

You can see which prefixes exist by issuing the z/OS command DISPLAY OPDATA.

If you are running in a sysplex, z/OS diagnoses any conflicts of this type at the time of CPF registration (see “CPFs in a sysplex environment” on page 1208 for information about CPF registration).

Table 149 on page 1208 shows the characters that you can use when defining your CPF strings:

Table 149. Valid character set for CPF strings

Character set	Contents
Alphabetic	Uppercase A through Z, lowercase a through z
Numeric	0 through 9
National (see note)	@ \$ # (Characters that can be represented as hexadecimal values)
Special	. ^ ( ) * & + - = < ! ! ; % _ ? : >

**Note:**

The system recognizes the following hexadecimal representations of the national characters: @ as X'7C', \$ as X'5B', and # as X'7B'. In countries other than the U.S., the U.S. national characters represented on terminal keyboards might generate a different hexadecimal representation and cause an error. For example, in some countries the \$ character might generate an X'4A'.

The semicolon (;) is valid as a CPF but on most systems, this character is the command delimiter.

**CPFs in a sysplex environment:**

Use this topic to understand how to use CPFs within the scope of a sysplex.

If used in a sysplex environment, IBM MQ registers your CPFs to enable you to enter a command from any console in the sysplex and route that command to the appropriate system for execution. The command responses are returned to the originating console.

**Defining the scope for sysplex operation**

Scope is used to determine the type of CPF registration performed by the IBM MQ subsystem when you are running IBM MQ in a sysplex environment.

Possible values for scope are as follows:

**M** System scope.

The CPF is registered with z/OS at system IPL time by IBM MQ and remains registered for the entire time that the z/OS system is active.

IBM MQ commands must be entered at a console connected to the z/OS image running the target subsystem, or you must use ROUTE commands to direct the command to that image.

Use this option if you are not running in a sysplex.

**S** Sysplex started scope.

The CPF is registered with z/OS when the IBM MQ subsystem is started, and remains active until the IBM MQ subsystem terminates.

You must use ROUTE commands to direct the original START QMGR command to the target system, but all further IBM MQ commands can be entered at any console connected to the sysplex, and are routed to the target system automatically.

After IBM MQ termination, you must use the ROUTE commands to direct subsequent START commands to the target IBM MQ subsystem.

**X** Sysplex IPL scope.

The CPF is registered with z/OS at system IPL time by IBM MQ and remains registered for the entire time that the z/OS system is active.



IBM MQ commands can be entered at any console connected to the sysplex, and are routed to the image that is executing the target system automatically.

An IBM MQ subsystem with a CPF with scope of S can be defined on one or more z/OS images within a sysplex, so these images can share a single subsystem name table. However, you must ensure that the initial START command is issued on (or routed to) the z/OS image on which you want the IBM MQ subsystem to run. If you use this option, you can stop the IBM MQ subsystem and restart it on a different z/OS image within the sysplex without having to change the subsystem name table or perform an IPL of a z/OS system.

An IBM MQ subsystem with a CPF with scope of X can only be defined on one z/OS image within a sysplex. If you use this option, you must define a unique subsystem name table for each z/OS image requiring IBM MQ subsystems with CPFs of scope X.

If you want to use the z/OS automatic restart manager (ARM) to restart queue managers in different z/OS images automatically, every queue manager must be defined in each z/OS image on which that queue manager might be restarted. Every queue manager must be defined with a sysplex-wide, unique 4-character subsystem name with a CPF scope of S.

### **Task 6: Create procedures for the IBM MQ queue manager**

Each IBM MQ subsystem needs a cataloged procedure to start the queue manager. You can create your own or use the IBM-supplied procedure library.

- Repeat this task for each IBM MQ queue manager.
- You might need to modify the cataloged procedure when migrating from a previous version.

For each IBM MQ subsystem defined in the subsystem name table, create a cataloged procedure in a procedure library for starting the queue manager. The IBM-supplied procedure library is called SYS1.PROCLIB, but your installation might use its own naming convention.

The name of the queue manager started task procedure is formed by concatenating the subsystem name with the characters MSTR. For example, subsystem CSQ1 has the procedure name CSQ1MSTR. You need one procedure for each subsystem you define.

Many examples and instructions in this product documentation assume that you have a subsystem called CSQ1. You might find these examples easier to use if a subsystem called CSQ1 is created initially for installation verification and testing purposes.

Two sample started task procedures are provided in thlqual.SCSQPROC. Member CSQ4MSTR uses one page set for each class of message, member CSQ4MSRR uses multiple page sets for the major classes of message. Copy one of these procedures to member xxxxMSTR (where xxxx is the name of your IBM MQ subsystem) of your SYS1.PROCLIB or, if you are not using SYS1.PROCLIB, your procedure library. Copy the sample procedure to a member in your procedure library for each IBM MQ subsystem that you define.

When you have copied the members, you can tailor them to the requirements of each subsystem, using the instructions in the member. For information about specifying region sizes below the 16 MB line, above the 16 MB line, and above the 2 GB bar, see Suggested region sizes. You can also use symbolic parameters in the JCL to allow the procedure to be modified when it is started. If you have several IBM MQ subsystems, you might find it advantageous to use JCL include groups for the common parts of the procedure, to simplify future maintenance.

If you are using queue-sharing groups, the STEPLIB concatenation must include the Db2 runtime target library SDSNLOAD, and it must be APF-authorized. This library is only required in the STEPLIB concatenation if it is not accessible through the link list or LPA.

If you are using IBM MQ Advanced Message Security the STEPLIB concatenation must include *thlqual.SDRQAUTH* and it must be APF authorized.

You can add the exit library (CSQXLIB) to this procedure later if you want to use queue manager exits. You need access to the Language Environment (LE) runtime library SCEERUN to do this; if it is not in your link list (SYS1.PARMLIB(LNKLSTkk)), concatenate it in the STEPLIB DD statement. You also must stop and restart your queue manager.

**Note:** You can make a note of the names of your bootstrap data set (BSDS), logs, and page sets for use in JCL and then define these sets at a later step in the process.

**Related concepts:**

“Task 7: Create procedures for the channel initiator”

For each IBM MQ subsystem, tailor a copy of CSQ4CHIN. Depending on what other products you are using, you might need to allow access to other data sets.

**Task 7: Create procedures for the channel initiator**

For each IBM MQ subsystem, tailor a copy of CSQ4CHIN. Depending on what other products you are using, you might need to allow access to other data sets.

- Repeat this task for each IBM MQ queue manager.
- You might need to perform this task when migrating from a previous version.

You need to create a channel-initiator started task procedure for each IBM MQ subsystem that is going to use distributed queuing.

To do this:

1. Copy the sample started task procedure *thlqual.SCSQPROC(CSQ4CHIN)* to your procedure library. Name the procedure *xxxx CHIN*, where *xxxx* is the name of your IBM MQ subsystem (for example, *CSQ1CHIN* would be the channel initiator started task procedure for queue manager *CSQ1*).
2. Make a copy for each IBM MQ subsystem that you are going to use.
3. Tailor the procedures to your requirements using the instructions in the sample procedure *CSQ4CHIN*. You can also use symbolic parameters in the JCL to allow the procedure to be modified when it is started. This is described with the start options in the *../com.ibm.mq.adm.doc/q022070\_.dita*.

Concatenate the distributed queuing library *thlqual.SCSQMVR1*.

Access to the LE runtime library *SCEERUN* is required; if it is not in your link list (SYS1.PARMLIB(LNKLSTkk)), concatenate it in the STEPLIB DD statement.

4. Authorize the procedures to run under your external security manager.

The channel initiator is a long running address space. To prevent its termination after a restricted amount of CPU has been consumed, confirm that either:

- The default for started tasks in your z/OS system is unlimited CPU; a JES2 configuration statement for *JOBCLASS(STC)* with *TIME=(1440,00)* achieves this, or
- Explicitly add a *TIME=1440*, or *TIME=NOLIMIT*, parameter to the EXEC statement for *CSQXJST*.

You can add the exit library (CSQXLIB) to this procedure later if you want to use channel exits. You need to stop and restart your channel initiator to do this.

If you are using SSL, access to the system Secure Sockets Layer (SSL) runtime library is required. This library is called *SIEALNKE*. The library must be APF authorized.

If you are using TCP/IP, the channel initiator address space must be able to access the TCPIP.DATA data set that contains TCP/IP system parameters. The ways that the data set has to be set up depends on which TCP/IP product and interface you are using. They include:

- Environment variable, RESOLVER\_CONFIG
- HFS file, /etc/resolv.conf
- //SYSTCPD DD statement
- //SYSTCPDD DD statement
- *jobname/userid*.TCPIP.DATA
- SYS1.TCPPARMS(TCPDATA)
- *zapname*.TCPIP.DATA

Some of these affect your started-task procedure JCL. For more information, see *z/OS Communications Server: IP Configuration Guide*.

**Related concepts:**

“Task 8: Define the IBM MQ subsystem to a z/OS WLM service class”

To give IBM MQ appropriate performance priority in the z/OS system, you must assign the queue manager and channel initiator address spaces to an appropriate z/OS workload management (WLM) service class. If you do not do this explicitly, inappropriate defaults might apply.

**Task 8: Define the IBM MQ subsystem to a z/OS WLM service class**

To give IBM MQ appropriate performance priority in the z/OS system, you must assign the queue manager and channel initiator address spaces to an appropriate z/OS workload management (WLM) service class. If you do not do this explicitly, inappropriate defaults might apply.

- *Repeat this task for each IBM MQ queue manager.*
- *You do not need to perform this task when migrating from a previous version.*

Use the ISPF dialog supplied with WLM to perform the following tasks:

- Extract the z/OS WLM policy definition from the WLM couple data set.
- Update this policy definition by adding queue manager and channel initiator started task procedure names to the chosen service class
- Install the changed policy on the WLM couple data set

Then activate this policy using the z/OS command

```
V WLM,POLICY=poli c yname,REFRESH
```

See for more information on setting performance options.

**Related concepts:**

“Task 9: Set up the Db2 environment”

If you are using queue-sharing groups you must create the required Db2 objects by customizing and running a number of sample jobs.

**Task 9: Set up the Db2 environment**

If you are using queue-sharing groups you must create the required Db2 objects by customizing and running a number of sample jobs.

For more information about selecting which offload storage environment to use, see *Deciding your offload storage environment*.

If you choose Db2 as the offload storage environment, see “Set up the Db2 environment” on page 1212. If you choose SMDS as the offload storage environment, see *Set up the SMDS environment*. If you are choosing the SMDS offload storage environment, you are still required to set up the Db2 environment for shared queues.

## Set up the Db2 environment

You must create and bind the required Db2 objects by customizing and running a number of sample jobs.

- Repeat this task for each Db2 data-sharing group.
- You might need to perform this task when migrating from a previous version.
- Omit this task if you are not using queue-sharing groups.

If you later want to use queue-sharing groups, perform this task at that time.

You must establish an environment in which IBM MQ can access and execute the Db2 plans that are used for queue-sharing groups.

The following steps must be performed for each new Db2 data-sharing group. All the sample JCL is in thlqual.SCSQPROC.

1. Customize and execute sample JCL CSQ45CSG to create the storage group that is to be used for the IBM MQ database, table spaces, and tables.
2. Customize and execute sample JCL CSQ45CDB to create the database to be used by all queue managers that are connecting to this Db2 data-sharing group.
3. Customize and execute sample JCL CSQ45CTS to create the table spaces that contain the queue manager and channel initiator tables used for queue-sharing groups (to be created in step 1 ).
4. Customize and execute sample JCL CSQ45CTB to create the 12 Db2 tables and associated indexes. Do not change any of the row names or attributes.
5. Customize and execute sample JCL CSQ45BPL to bind the Db2 plans for the queue manager, utilities, and channel initiator.
6. Customize and execute sample JCL CSQ45GEX to grant execute authority to the plans for the user IDs that are used by the queue manager, utilities, and channel initiator. The user IDs for the queue manager and channel initiator are the user IDs under which their started task procedures run. The user IDs for the utilities are the user IDs under which the batch jobs can be submitted. The names of the appropriate plans are:

User	Plans
Queue manager	CSQ5A 800, CSQ5C 800, CSQ5D 800, CSQ5K 800, CSQ5L 800, CSQ5M 800, CSQ5P 800, CSQ5R 800, CSQ5S 800, CSQ5T 800, CSQ5U 800, CSQ5W 800
SDEFS function of the CSQUTIL batch utility	CSQ52 800
CSQ5PQSG and CSQJUCNV batch utilities	CSQ5B 800
CSQUZAP service utility	CSQ5Z 800

In the event of a failure during Db2 setup, the following jobs can be customized and executed:

- CSQ45DTB to drop the tables and indexes.
- CSQ45DTS to drop the table spaces.
- CSQ45DDB to drop the database.
- CSQ45DSG to drop the storage group.

**Note:** If these jobs fail because of a Db2 locking problem it is probably due to contention for a Db2 resource, especially if the system is being heavily used. Resubmit the jobs later. It is preferable to run these jobs when the system is lightly used or quiesced.

See *Db2 10 for z/OS: Db2 Administration* for more information about setting up Db2.

See Planning on z/OS for information about Db2 table sizes.

## Set up the SMDS environment

You should use SMDS to offload large messages. You can offload all messages to SMDS to give you more capacity in your structure.

You can use Storage Class Memory(SCM); see Use of storage class memory with shared queues.

- Estimate structure and data set space requirements. See Shared message data set capacity considerations.
- Allocate and preformat data sets. See Creating a shared message data set.
- Use the following MQSC command to display the **CFLEVEL** and **OFFLOAD** status.

```
DISPLAY CFSTRUCT(*) CFLEVEL OFFLOAD
```

For more information about the **DISPLAY CFSTRUCT** command, see DISPLAY CFSTRUCT.

- Ensure that the coupling facility structure is defined with **CFLEVEL(5)** and **OFFLOAD(SMDS)** by using the following MQSC commands:

```
ALTER CFSTRUCT(APP1) CFLEVEL(5)
```

```
ALTER CFSTRUCT(APP1) OFFLOAD(SMDS)
```

For more information about the **ALTER CFSTRUCT** command, see ALTER CFSTRUCT.

### Related concepts:

“Task 10: Set up the coupling facility”

If you are using queue-sharing groups, define the coupling facility structures used by the queue managers in the queue-sharing group in the coupling facility Resource Management (CFRM) policy data set, using IXCMIAPU.

### Task 10: Set up the coupling facility

If you are using queue-sharing groups, define the coupling facility structures used by the queue managers in the queue-sharing group in the coupling facility Resource Management (CFRM) policy data set, using IXCMIAPU.

- *Repeat this task for each queue-sharing group.*
- *You might need to perform this task when migrating from a previous version.*
- *Omit this task if you are not using queue-sharing groups.*

*If you later want to use queue-sharing groups, perform this task at that time.*

All the structures for the queue-sharing group start with the name of the queue-sharing group. Define the following structures:

- An administrative structure called *qsg-name* CSQ\_ADMIN. This structure is used by IBM MQ itself and does not contain any user data.
- A system application structure called *qsg-name* CSQSYSAPPL. This structure is used by IBM MQ system queues to store state information.
- One or more structures used to hold messages for shared queues. These can have any name you choose up to 16 characters long.
  - The first four characters must be the queue-sharing group name. (If the queue-sharing group name is less than four characters long, it must be padded to four characters with @ symbols.)
  - The fifth character must be alphabetic and subsequent characters can be alphabetic or numeric. This part of the name (without the queue-sharing group name) is what you specify for the CFSTRUCT name when you define a shared queue, or a CF structure object.

You can use only alphabetic and numeric characters in the names of structures used to hold messages for shared queues, you cannot use any other characters (for example, the \_ character, which is used in the name of the administrative structure).

Sample control statements for IXCMIAPU are in data set thlqual.SCSQPROC(CSQ4CFRM). Customize these and add them to your IXCMIAPU job for the coupling facility and run it.

When you have defined your structures successfully, activate the CFRM policy that is being used. To do this, issue the following z/OS command:

```
SETXCF START,POLICY,TYPE=CFRM,POLNAME= policy-name
```

See the Defining coupling facility resources for information about planning CF structures and their sizes.

#### **Related concepts:**

“Task 11: Implement your ESM security controls”

Implement security controls for queue-sharing groups, the channel initiator, and all queue managers accessing the coupling facility list structures.

### **Task 11: Implement your ESM security controls**

Implement security controls for queue-sharing groups, the channel initiator, and all queue managers accessing the coupling facility list structures.

- Repeat this task for each IBM MQ queue manager or queue-sharing group.
- You might need to perform this task when migrating from a previous version.

If you use RACF as your external security manager, see Setting up security on z/OS , which describes how to implement these security controls.

If you are using queue-sharing groups, ensure that the user IDs associated with the queue manager, channel initiator, and the utilities (as specified in task 9, step 6 on page 1212 ) have authority to establish an RRSAF connection to each Db2 subsystem with which you want to establish a connection. The RACF profile to which the user ID requires READ access is *DB2ssid.RRSAF* in the DSNR resource class.

If you are using the channel initiator, you must also do the following:

- If your subsystem has connection security active, define a connection security profile *ssid.CHIN* to your external security manager (see Connection security profiles for the channel initiator for information about this).
- If you are using the Secure Sockets Layer (SSL) or a sockets interface, ensure that the user ID under whose authority the channel initiator is running is configured to use UNIX System Services, as described in the *OS/390® UNIX System Services Planning* documentation.
- If you are using SSL, ensure that the user ID under whose authority the channel initiator is running is configured to access the key ring specified in the *SSLKEYR* parameter of the ALTER QMGR command.

Those queue managers that will access the coupling facility list structures require the appropriate security access. The RACF class is FACILITY. The queue manager user ID requires ALTER access to the *IXLSTR. structure-name* profile.

Before you start the queue manager, set up IBM MQ data set and system security by:

- Authorizing the queue manager started task procedure to run under your external security manager.
- Authorizing access to the queue manager data sets.

For details about how to do this, see Security installation tasks for z/OS(r).

If you are using RACF, provided you use the RACF STARTED class, you do not need to perform an IPL of your system (see RACF authorization of started-task procedures ).

**Related concepts:**

“Task 12: Update SYS1.PARMLIB members”

To ensure that your changes remain in effect after an IPL, you must update some members of SYS1.PARMLIB

**Task 12: Update SYS1.PARMLIB members**

To ensure that your changes remain in effect after an IPL, you must update some members of SYS1.PARMLIB

- *You need to perform this task once for each z/OS system where you want to run IBM MQ.*
- *If you are using queue-sharing groups, you must ensure that the settings for IBM MQ are identical on each z/OS system in the sysplex.*
- *You might need to perform this task when migrating from a previous version.*

Update SYS1.PARMLIB members as follows:

1. Update member IEFSSNss as described in “Task 5: Define the IBM MQ subsystem to z/OS” on page 1205.
2. Change IEASYSpP so that the following members are used when an IPL is performed:
  - the PROGxx or IEAAPFaa members used in “Task 2: APF authorize the IBM MQ load libraries” on page 1202
  - the LNKLSTkk and LPALSTmm members used in “Task 3: Update the z/OS link list and LPA” on page 1203
  - the SCHEDxx member used in “Task 4: Update the z/OS program properties table” on page 1205
  - the IEFSSNss member used in “Task 5: Define the IBM MQ subsystem to z/OS” on page 1205

**Related concepts:**

“Task 13: Customize the initialization input data sets”

Make working copies of the sample initialization input data sets and tailor them to suit your system requirements.

**Task 13: Customize the initialization input data sets**

Make working copies of the sample initialization input data sets and tailor them to suit your system requirements.

- *Repeat this task for each IBM MQ queue manager.*
- *You need to perform this task when migrating from a previous version.*

Each IBM MQ queue manager gets its initial definitions from a series of commands contained in the IBM MQ *initialization input data sets*. These data sets are referenced by the DDnames CSQINP1, CSQINP2 and CSQINPT defined in the queue manager started task procedure.

Responses to these commands are written to the initialization output data sets referenced by the DDnames CSQOUT1, CSQOUT2 and CSQOUTT.

To preserve the originals, make working copies of each sample. Then you can tailor the commands in these working copies to suit your system requirements.

If you use more than one IBM MQ subsystem, if you include the subsystem name in the high-level qualifier of the initialization input data set name, you can identify the IBM MQ subsystem associated with each data set more easily.

Refer to the following topics for further information about the samples:

- Initialization data set formats
- Using the CSQINP1 sample
- Using the CSQINP2 samples

- Using the CSQINPX sample
- Using the CSQINPT sample

### Initialization data set formats

The initialization input data sets can be partitioned data set (PDS) members or sequential data sets. They can be a concatenated series of data sets. Define them with a record length of 80 bytes, where:

- Only columns 1 through 72 are significant. Columns 73 through 80 are ignored.
- Records with an asterisk (\*) in column 1 are interpreted as comments and are ignored.
- Blank records are ignored.
- Each command must start on a new record.
- A trailing - means continue from column 1 of the next record.
- A trailing + means continue from the first non-blank column of the next record.
- The maximum number of characters permitted in a command is 32 762.

The initialization output data sets are sequential data sets, with a record length of 125, a record format of VBA, and a block size of 629.

### Using the CSQINP1 sample

Data set thlqual.SCSQPROC holds two members which contain definitions of buffer pools, page set to buffer pool associations, and an ALTER SECURITY command.

Member CSQ4INP1 uses one page set for each class of message. Member CSQ4INPR uses multiple page sets for the major classes of message.

Include the appropriate sample in the CSQINP1 concatenation of your queue manager started task procedure.

#### Notes:

1. IBM MQ supports up to 16 buffer pools (zero through 15). If OPMODE is set to OPMODE=(NEWFUNC, 800), 100 buffer pools are supported in the range zero through 99. The DEFINE BUFFPOOL command can only be issued from a CSQINP1 initialization data set. The definitions in the sample specify four buffer pools.
2. Each page set used by the queue manager must be defined in the CSQINP1 initialization data set by using the DEFINE PSID command. The page set definition associates a buffer pool ID with a page set. If no buffer pool is specified, buffer pool zero is used by default.  
Page set zero (00) must be defined. It contains all the object definitions. You can define up to 100 page sets for each queue manager.
3. The ALTER SECURITY command can be used to alter the security attributes TIMEOUT and INTERVAL. In CSQ4INP1, the default values are defined as 54 for TIMEOUT and 12 for INTERVAL.

See the Planning on z/OS for information about organizing buffer pools and page sets.

If you change the buffer pool and page set definitions dynamically while the queue manager is running, you should also update the CSQINP1 definitions. The changes are only retained for a cold start of IBM MQ, unless the buffer pool definition includes the REPLACE attribute.

### Using the CSQINP2 samples

This table lists the members of thlqual.SCSQPROC that can be included in the CSQINP2 concatenation of your queue manager started task procedure, with a description of their function. The naming convention is CSQ4INS\*. CSQ4INY\* will need to be modified for YOUR configuration. You should avoid changing



CSQINS\* members because you will need to reapply any changes when you migrate to the next release. Instead, you can put DEFINE or ALTER commands in CSQ4INY\* members.

Table 150. Members of thlqual.SCSQPROC

Member name	Description
CSQ4INSG	System object definitions.
CSQ4INSA	System object and default rules for channel authentication.
CSQ4INSX	System object definitions.
CSQ4INSS	Customize and include this member if you are using queue-sharing groups.
CSQ4INSJ	Customize and include this member if you are using publish/subscribe using JMS.
CSQ4INSM	System object definitions for advanced message security.
CSQ4INSR	Customize and include this member if you are using WebSphere Application Server, or the queued publish/subscribe interface supported by the queued publish/subscribe daemon in IBM MQ V7 or later.
CSQ4DISP	CSQINP2 sample for displaying object definitions.
CSQ4INYC	Clustering definitions.
CSQ4INYD	Distributed queuing definitions.
CSQ4INYG	General definitions.
CSQ4INYR	Storage class definitions, using multiple page sets for the major classes of message.
CSQ4INYS	Storage class definitions, using one page set for each class of message.

You need to define objects once only, not each time that you start a queue manager, so it is not necessary to include these definitions in CSQINP2 every time. If you do include them every time, you are attempting to define objects that already exist, and you will get messages similar to the following:

```
CSQM095I +CSQ1 CSQMAQLC QLOCAL(SYSTEM.DEFAULT.LOCAL.QUEUE) ALREADY EXISTS
CSQM090E +CSQ1 CSQMAQLC FAILURE REASON CODE X'00D44003'
CSQ9023E +CSQ1 CSQMAQLC ' DEFINE QLOCAL' ABNORMAL COMPLETION
```

The objects are not damaged by this failure. If you want to leave the SYSTEM definitions data set in the CSQINP2 concatenation, you can avoid the failure messages by specifying the REPLACE attribute against each object.

### Using the CSQINPX sample

Sample thlqual.SCSQPROC(CSQ4INPX) contains a set of commands that you might want to execute each time the channel initiator starts. These are typically channel-related commands such as START LISTENER, which are required every time the channel initiator starts, rather than whenever the queue manager starts, and which are not allowed in the input data sets CSQINP1 or CSQINP2. You must customize this sample before use; you can then include it in the CSQINPX data set for the channel initiator.

The IBM MQ commands contained in the data set are executed at the end of channel initiator initialization, and output is written to the data set specified by the CSQOUTX DD statement. The output is like that produced by the COMMAND function of the IBM MQ utility program (CSQUTIL). See The CSQUTIL utility for more details.

You can specify any of the IBM MQ commands that can be issued from CSQUTIL, not only the channel commands. You can enter commands from other sources while CSQINPX is being processed. All commands are issued in sequence, regardless of the success of the previous command.

To specify a command response time, you can use the pseudo-command `COMMAND` as the first command in the data set. This takes a single optional keyword `RESPTIME( mmm )`, where *mmm* is the time, in seconds, to wait for a response to each command. This is in the range 5 through 999; the default is 30.

If IBM MQ detects that the responses to four commands have taken too long, processing of `CSQINPX` is stopped and no further commands are issued. The channel initiator is not stopped, but message `CSQU052E` is written to the `CSQOUTX` data set, and message `CSQU013E` is sent to the console.

When IBM MQ has completed processing of `CSQINPX` successfully, message `CSQU012I` is sent to the console.

## Using the `CSQINPT` sample

This table lists the members of `thlqual.SCSQPROC` that can be included in the `CSQINPT` concatenation of your queue manager started task procedure, with a description of their function.

*Table 151. Members of thlqual.SCSQPROC*

Member name	Description
<code>CSQ4INST</code>	System default subscription definition.
<code>CSQ4INYT</code>	Publish/Subscribe definitions.

The IBM MQ commands contained in the data set are executed when publish/subscribe initialization completes, and output is written to the data set specified by the `CSQOUTT DD` statement. The output is like that produced by the `COMMAND` function of the IBM MQ utility program (`CSQUTIL`). See The `CSQUTIL` utility for more details.

### Related concepts:

“Task 14: Create the bootstrap and log data sets”

Use the supplied program `CSQJU003` to prepare the bootstrap data sets (BSDSs) and log data sets.

### Task 14: Create the bootstrap and log data sets

Use the supplied program `CSQJU003` to prepare the bootstrap data sets (BSDSs) and log data sets.

- Repeat this task for each IBM MQ queue manager.
- You do not need to perform this task when migrating from a previous version.

The sample JCL and Access Method Services (AMS) control statements to run `CSQJU003` to create a single or dual logging environment are held in `thlqual.SCSQPROC(CSQ4BSDS)`. Customize and run this job to create your BSDSs and logs and to preformat the logs.

**Important:** You should use the newest version of `CSQ4BSDS`, or update your JCL manually to use `RECORDS(850 60)`.

The started task procedure, `CSQ4MSTR`, described in “Task 6: Create procedures for the IBM MQ queue manager” on page 1209, refers to BSDSs in statements of the form:

```
//BSDS1 DD DSN=++HLQ++.BSDS01,DISP=SHR
//BSDS2 DD DSN=++HLQ++.BSDS02,DISP=SHR
```

The log data sets are referred to by the BSDSs.

**Note:**

1. The BLKSIZE must be specified on the SYSPRINT DD statement in the LOGDEF step. The BLKSIZE must be 629.
2. To help identify bootstrap data sets and log data sets from different queue managers, include the subsystem name in the high level qualifier of these data sets.
3. If you are using queue-sharing groups, you must define the bootstrap and log data sets with SHAREOPTIONS(2 3).

See Planning on z/OS for information about planning bootstrap and log data sets and their sizes.

For IBM MQ Version 8.0, the 8 byte log RBA enhancement improves the availability of a queue manager, as described in Larger log Relative Byte Address. To enable 8 byte log RBA on a queue manager before the queue manager is first started, perform the following steps after creating your logging environment.

1. Using **IDCAMS ALTER**, rename the version 1 format BSDSs (created using the CSQJU003 program) to something like ++HLQ++.V1.BSDS01.

**Note:** Ensure that you rename the data and index components as well as the VSAM cluster.

2. Allocate new BSDSs with the same attributes as the ones already defined. These will become the version 2 format BSDSs that will be used by the queue manager when it is started.
3. Run the BSDS conversion utility (CSQJUCNV) to convert the version 1 format BSDSs to the new version 2 format BSDSs.
4. Once the conversion completes successfully, delete the version 1 format BSDSs.
5. In order to use 8 byte log RBA, ensure that Version 8.0 new functions are enabled with OPMODE as described in “Task 17: Tailor your system parameter module” on page 1221.

**Note:** If the queue manager is in a queue-sharing group, all queue managers in the queue-sharing group must have been started with OPMODE(NEWFUNC,800) before 8 byte log RBA can be enabled.

**Related concepts:**

“Task 15: Define your page sets”

Define page sets for each queue manager using one of the supplied samples.

### **Task 15: Define your page sets**

Define page sets for each queue manager using one of the supplied samples.

- Repeat this task for each IBM MQ queue manager.
- You do not need to perform this task when migrating from a previous version.

Define separate page sets for each IBM MQ queue manager. thlqual.SCSQPROC(CSQ4PAGE) and thlqual.SCSQPROC(CSQ4PAGR) contain JCL and AMS control statements to define and format page sets. Member CSQ4PAGE uses one page set for each class of message, member CSQ4PAGR uses multiple page sets for the major classes of message. The JCL runs the supplied utility program CSQUTIL. Review the samples and customize them for the number of page sets you want and the sizes to use. See the Planning on z/OS for information about page sets and how to calculate suitable sizes.

The started task procedure CSQ4MSTR described in “Task 6: Create procedures for the IBM MQ queue manager” on page 1209 refers to the page sets, in a statement of the form:

```
//CSQP00 nn DD DISP=OLD,DSN= xxxxxxxxx
```

where *nn* is the page set number between 00 and 99, and *xxxxxxx* is the data set that you define.

**Note:**

1. If you intend to use the dynamic page set expansion feature, ensure that secondary extents are defined for each page set. `thlqual.SCSQPROC(CSQ4PAGE)` shows how to do this.
2. To help identify page sets from different queue managers, include the subsystem name in the high level qualifier of the data set associated with each page set.
3. If you intend to allow the FORCE option to be used with the FORMAT function of the utility program CSQUTIL, you must add the REUSE attribute on the AMS DEFINE CLUSTER statement. This is described in the Administering IBM MQ for z/OS.
4. If your page sets are to be larger than 4 GB you must use the Storage Management System (SMS) EXTENDED ADDRESSABILITY function.

**Related concepts:**

*"Task 16: Add the IBM MQ entries to the Db2 tables"*

If you are using queue-sharing groups, run the CSQ5PQSG utility to add queue-sharing group and queue manager entries to the IBM MQ tables in the Db2 data-sharing group.

**Task 16: Add the IBM MQ entries to the Db2 tables**

If you are using queue-sharing groups, run the CSQ5PQSG utility to add queue-sharing group and queue manager entries to the IBM MQ tables in the Db2 data-sharing group.

- *Repeat this task for each IBM MQ queue-sharing group and each queue manager.*
- *You might need to perform this task when migrating from a previous version.*
- *Omit this task if you are not using queue-sharing groups.*

*If you later want to use queue-sharing groups, perform this task at that time.*

Run CSQ5PQSG for each queue-sharing group and each queue manager that is to be a member of a queue-sharing group. (CSQ5PQSG is described in the Administering IBM MQ for z/OS.)

Perform the following actions in the specified order:

1. Add a queue-sharing group entry into the IBM MQ Db2 tables using the ADD QSG function of the CSQ5PQSG program. A sample is provided in `thlqual.SCSQPROC(CSQ45AQS)`.  
Perform this function once for each queue-sharing group that is defined in the Db2 data-sharing group. The queue-sharing group entry must exist before adding any queue manager entries that reference the queue-sharing group.
2. Add a queue manager entry into the IBM MQ Db2 tables using the ADD QMGR function of the CSQ5PQSG program. A sample is provided in `thlqual.SCSQPROC(CSQ45AQM)`.  
Perform this function for each queue manager that is to be a member of the queue-sharing group.

**Note:**

- a. A queue manager can only be a member of one queue-sharing group.
- b. You must have RRS running to be able to use queue-sharing groups.

## Related concepts:

“Task 17: Tailor your system parameter module”

The IBM MQ system parameter module controls the logging, archiving, tracing, and connection environments that IBM MQ uses in its operation. A default module is supplied, or you can create your own using supplied JCL and assembler source modules.

## Task 17: Tailor your system parameter module

The IBM MQ system parameter module controls the logging, archiving, tracing, and connection environments that IBM MQ uses in its operation. A default module is supplied, or you can create your own using supplied JCL and assembler source modules.

- Repeat this task for each IBM MQ queue manager, as required.
- You need to perform this task when migrating from a previous version. For details, see *Migrating and upgrading IBM MQ*.
- To enable IBM MQ Advanced Message Security for z/OS on an existing queue manager, you only need to set *SPLCAP* to *YES* as described in “Using *CSQ6SYSP*” on page 1223. If you are configuring this queue manager for the first time, complete the whole of this task.

The system parameter module has three macros as follows:

Macro name	Purpose
CSQ6SYSP	Specifies the connection and tracing parameters, see “Using <i>CSQ6SYSP</i> ” on page 1223
CSQ6LOGP	Controls log initialization, see “Using <i>CSQ6LOGP</i> ” on page 1232
CSQ6ARVP	Controls archive initialization, see “Using <i>CSQ6ARVP</i> ” on page 1235

IBM MQ supplies a default system parameter module, *CSQZPARM*, which is invoked automatically if you issue the *START QMGR* command (without a *PARM* parameter) to start an instance of IBM MQ. *CSQZPARM* is in the APF-authorized library *thlqual.SCSQAUTH* also supplied with IBM MQ. The values of these parameters are displayed as a series of messages when you start IBM MQ.

See *START QMGR* for more information about how this command is used.

## Creating your own system parameter module

If *CSQZPARM* does not contain the system parameters you want, you can create your own system parameter module using the sample JCL provided in *thlqual.SCSQPROC(CSQ4ZPRM)*.

To create your own system parameter module:

1. Make a working copy of the JCL sample.
2. Edit the parameters for each macro in the copy as required. If you remove any parameters from the macro calls, the default values are automatically picked up at run time.
3. Replace the placeholder *++NAME++* with the name that the load module is to take (this can be *CSQZPARM*).
4. If your assembler is not high-level assembler, change the JCL as required by your assembler.
5. Run the JCL to assemble and link edit the tailored versions of the system parameter macros to produce a load module. This is the new system parameter module with the name that you have specified.
6. Put the load module produced in an APF-authorized user library.
7. Add user *READ* access to the APF-authorized user library.
8. Include this library in the IBM MQ queue manager started task procedure *STEPLIB*. This library name must come before the library *thlqual.SCSQAUTH* in *STEPLIB*.

- Invoke the new system parameter module when you start the queue manager. For example, if the new module is named NEWMODS, issue the command:

```
START QMGR PARM(NEWMODS)
```

- Ensure successful completion of the command by checking the job log. There should be an entry in the log similar to the following:

```
CSQ9022I CDL1 CSQYASCP 'START QMGR' NORMAL COMPLETION
```

You can also specify the parameter module name in the queue manager startup JCL. For more information, see *Starting and stopping a queue manager*.

**Note:** If you choose to name your module CSQZPARM, you do not need to specify the PARM parameter on the START QMGR command.

### Fine tuning a system parameter module

IBM MQ also supplies a set of three assembler source modules, which can be used to fine-tune an existing system parameter module. These modules are in library thlqual.SCSQASMS. Typically, you use these modules in a test environment to change the default parameters in the system parameter macros. Each source module calls a different system parameter macro:

This assembler source module...	Calls this macro...
CSQFSYSP	CSQ6SYSP (connection and tracing parameters)
CSQJLOGP	CSQ6LOGP (log initialization)
CSQJARVP	CSQ6ARVP (archive initialization)

This is how you use these modules:

- Make working copies of each assembler source module in a user assembler library.
- Edit your copies by adding or altering the values of any parameters as required.
- Assemble your copies of any edited modules to create object modules in a user object library.
- Link edit these object code modules with an existing system parameter module to produce a load module that is the new system parameter module.
- Ensure that new system parameter module is a member of a user authorized library.
- Include this library in the queue manager started task procedure STEPLIB. This library must come before the library thlqual.SCSQAUTH in STEPLIB.
- Invoke the new system parameter module by issuing a START QMGR command, specifying the new module name in the PARM parameter, as before.

A sample usermod is provided in member CSQ4UZPR of SCSQPROC which demonstrates how to manage customized system parameters under SMP/E control.

### Altering system parameters

You can alter some system parameters while a queue manager is running; see the SET SYSTEM, SET LOG, and SET ARCHIVE commands.

Put the SET commands in your initialization input data sets so that they take effect every time you start the queue manager.

**Related concepts:**

“Task 18: Tailor the channel initiator parameters” on page 1242

Use ALTER QMGR to customize the channel initiator to suit your requirements.

**Using CSQ6SYSP:**

Use this topic as a reference for how to set system parameters using CSQ6SYSP.

The default parameters for CSQ6SYSP, and whether you can alter each parameter using the SET SYSTEM command, are shown in Table 152. If you want to change any of these values, see the detailed descriptions of the parameters.

Table 152. Default values of CSQ6SYSP parameters

Parameter	Description	Default value	SET command
ACELIM	Size of ACE storage pool in 1 KB blocks.	0 (no limit)	✓
CLCACHE	Specifies the type of cluster cache to use.	STATIC	-
CMDUSER	The default user ID for command security checks.	CSQOPR	-
CONNSWAP	Specifies whether jobs that are issuing certain IBM MQ API calls are swappable or non-swappable.	YES	-
EXCLMSG	Specifies a list of messages to be excluded from any log. Messages in this list are not sent to the z/OS console and hardcopy log. As a result using the EXCLMSG parameter to exclude messages is more efficient from a CPU perspective than using the methods described in “Task 22: Suppress information messages” on page 1247	( )	✓
EXITLIM	Time (in seconds) for which queue-manager exits can run during each invocation.	30	-
EXITTCB	How many started server tasks to use to run queue manager exits.	8	-
LOGLOAD	Number of log records written by IBM MQ between the start of one checkpoint and the next.	500 000	✓
MULCCAPT	Determines the Measured Usage Pricing property which controls the algorithm for gathering data used by Measured Usage License Charging (MULC).	See parameter description	-
OPMODE	Controls the operation mode of the queue manager.	See parameter description	-
OTMACON	OTMA connection parameters.	See parameter description	-
QINDEXBLD	Determines whether queue manager restart waits until all indexes are rebuilt, or completes before all indexes are rebuilt.	WAIT	-
QMCCSID	Coded character set identifier for the queue manager.	Zero	-
QSGDATA	Queue-sharing group parameters.	See parameter description	-
RESAUDIT	RESLEVEL auditing parameter.	YES	-
ROUTCDE	Message routing code assigned to messages not solicited from a specific console.	1	-
SERVICE	Reserved for use by IBM.	0	✓

Table 152. Default values of CSQ6SYSP parameters (continued)

Parameter	Description	Default value	SET command
SMFACCT	Specifies whether SMF accounting data is to be collected when the queue manager is started.  Note that class 4 channel accounting data is collected only when the channel initiator is started.	NO	-
SMFSTAT	Specifies whether SMF statistics are to be collected when the queue manager is started.  Note that class 4 channel initiator statistics data is collected only when the channel initiator is started.	NO	-
SPLCAP	Specifies whether queue security policy capability is enabled on this queue manager. For IBM MQ Advanced Message Security for z/OS, set this parameter to YES.	NO	-
STATIME	Default time, in minutes, between each gathering of statistics.	30	✓
TRACSTR	Specifies whether tracing is to be started automatically.	NO	-
TRACTBL	Size of trace table, in 4 KB blocks, to be used by the global trace facility.	99 (396 KB)	✓
WLMTIME	Time between scanning the queue index for WLM-managed queues.	30	-
WLMTIMU	Units (minutes or seconds) for WLMTIME.	MINS	-

#### ACELIM

Specifies the maximum size of the ACE storage pool in 1 KB blocks. The number must be in the range 0-999999. The default value of zero means no imposed constraint, beyond what is available in the system.

You should only set a value for ACELIM on queue managers that have been identified as using exorbitant quantities of ECSA storage. Limiting the ACE storage pool has the effect of limiting the number of connections in the system, and so, the amount of ECSA storage used by a queue manager.

Once the queue manager reaches the limit it is not possible for applications to obtain new connections. The lack of new connections causes failures in MQCONN processing, and applications coordinated through RRS are likely to experience failures in any IBM MQ API.

An ACE represents approximately 8% of the total ECSA required for the thread-related control blocks for a connection. So, for example, specifying ACELIM=5120 would be expected to cap the total amount of ECSA allocated by the queue manager (for thread-related control blocks) at approximately 64000K; that is 5120 multiplied by 12.5.

In order to cap the amount total amount of ECSA allocated by the queue-manager, for thread-related control blocks at 5120K, an ACELIM value of 410 is required.

You can use SMF 115 subtype 5 records, produced by statistics CLASS(3) trace, to monitor the size of the 'ACE/PEB' storage pool, and hence set an appropriate value for ACELIM.

You can obtain the total amount of ECSA storage used by the queue-manager, for control blocks, from SMF 115 subtype 7 records, written by statistics CLASS(2) trace; that is the first two elements in QSRSPHBT added together.

Note that, you should consider setting ACELIM as a mechanism to protect a z/OS image from a badly behaving queue manager, rather than as a means to control application connections to a queue manager.



**CLCACHE**

Specifies the type of cluster cache to use. See “Configuring a queue manager cluster” on page 902 for more information.

**STATIC**

When the cluster cache is static, its size is fixed at queue manager start-up, enough for the current amount of cluster information plus some space for expansion. The size cannot increase while the queue manager is active. This is the default.

**DYNAMIC**

When the cluster cache is dynamic, the initial size allocated at queue manager startup can be increased automatically if required while the queue manager is active.

**CMDUSER**

Specifies the default user ID used for command security checks. This user ID must be defined to the ESM (for example, RACF ). Specify a name of 1 through 8 alphanumeric characters. The first character must be alphabetic.

The default is CSQOPR.

**CONNSWAP**

Specifies whether batch jobs that are issuing certain IBM MQ API calls are swappable or non-swappable for the duration of the IBM MQ API request. Specify one of the following values:

**NO** Jobs are non-swappable during certain IBM MQ API calls.

**YES** Jobs are swappable during all IBM MQ API calls.

The default value is YES.

Use this parameter if low-priority jobs are swapped out while holding IBM MQ resources that other jobs or tasks might be waiting for. This parameter replaces the service parameter that was included IBM MQ V701; the service parameter is no longer in use.

IBM MQ views WebSphere Application Server as part of an RRSBATCH environment. From IBM WebSphere MQ Version 7.1, when the CONNSWAP keyword is used, it is applied to any application in a BATCH or RRSBATCH environment. The CONNSWAP keyword is also applicable to TSO users, however, it is not applicable for CICS or IMS applications. CONNSWAP changes are implemented when a recycle of the queue manager takes place. A recycle is required after the keyword change is made, because the CSQ6SYSP macro is reassembled, and the queue-manager restarted using the load module which is updated by the macro.

Alternatively, the WebSphere Application Server address space can be made non-swappable by using PPT.

**EXCLMSG**

Specifies a list of error messages to be excluded.

This list is dynamic and is updated using the SET SYSTEM command.

The default value is an empty list ( ).

Messages are supplied without the CSQ prefix and without the action code suffix (I-D-E-A). For example, to exclude message CSQX500I, add X500 to this list. This list can contain a maximum of 16 message identifiers.

To be eligible to be included in the list, the message must be issued after normal start up of the MSTR or CHIN address spaces and begin with the one of the following characters E, H, I, J, L, M, N, P, R, T, V, W, X, Y, 2, 3, 5, 9.

Message identifiers that are issued as a result of processing commands can be added to the list, however will not be excluded. For example, a message identifier is issued as a result of the DISPLAY USAGE PSID(\*) command, however, this message can not be suppressed.

## EXITLIM

Specifies the time, in seconds, allowed for each invocation of the queue manager exits. (This parameter has no effect on channel exits.)

Specify a value in the range 5 through 9999.

The default is 30. The queue manager polls exits that are running every 30 seconds. On each poll, any that have been running for more than the time specified by EXITLIM are forcibly terminated.

## EXITTCB

Specifies the number of started server tasks to use to run exits in the queue manager. (This parameter has no effect on channel exits.) You must specify a number at least as high as the maximum number of exits (other than channel exits) that the queue manager might have to run, else it will fail with a 6c6 abend.

Specify a value in the range zero through 99. A value of zero means that no exits can be run.

The default is 8.

## LOGLOAD

Specifies the number of log records that IBM MQ writes between the start of one checkpoint and the next. IBM MQ starts a new checkpoint after the number of records that you specify has been written.

Specify a value in the range 200 through 16 000 000.

The default is 500 000.

The greater the value, the better the performance of IBM MQ ; however, restart takes longer if the parameter is set to a large value.

Suggested settings:

<b>Test system</b>	10 000
<b>Production system</b>	500 000

In a production system, the supplied default value might result in a checkpoint frequency that is too high.

The value of LOGLOAD determines the frequency of queue manager checkpoints. Too large a value means that a large amount of data is written to the log between checkpoints, resulting in an increased queue manager forward recovery restart time following a failure. Too small a value causes checkpoints to occur too frequently during peak load, adversely affecting response times and processor usage.

An initial value of 500 000 is suggested for LOGLOAD. For a 1 KB persistent message rate of 100 messages a second (that is, 100 **MQPUT** s with commit and 100 **MQGET** s with commit) the interval between checkpoints is approximately 5 minutes.

**Note:** This is intended as a guideline only and the optimum value for this parameter is dependent on the characteristics of the individual system.

## MULCCAPT

Specifies the algorithm to be used for gathering data used by Measured Usage License Charging (MULC).

### STANDARD

MULC is based on the time from the IBM MQ API MQCONN call to the time of the IBM MQ API MQDISC call.

### REFINED

MULC is based on the time from the start of an IBM MQ API call to the end of the IBM MQ API call.

The default is STANDARD

**OPMODE**=( *Mode*,*VerificationLevel* )

OPMODE specifies the operation mode of the queue manager.

The default setting of OPMODE is OPMODE=(COMPAT,800) .

#### *Mode*

Specifies the requested operation mode. The values are as follows:

#### **COMPAT**

The queue manager runs in compatibility mode. Certain new functions are not available. The queue manager can be migrated back to an earlier release.

#### **NEWFUNC**

All new functions provided in this level of code are available. The queue manager cannot be migrated back to an earlier release.

#### *VerificationLevel*

*VerificationLevel* is a Version.Release.Modification (VRM) code, without punctuation; 800, for example.

The value of *VerificationLevel* ensures that the **CSQ6SYSP** parameters are coded for use with the level of **CSQ6SYSP** macro being compiled. If *VerificationLevel* does not match the VRM level of SCSQMACS used for **CSQ6SYSP**, then a compile-time error is reported. The *VerificationLevel* is compiled into the parameter module.

At queue manager startup, if the *VerificationLevel* does not match the release level of the queue manager, then COMPAT mode is forced.

The intent of the *VerificationLevel* parameter is to avoid inadvertent and irreversible setting of OPMODE to NEWFUNC. The mistake might occur when migrating to a newer version of IBM MQ using **CSQ6SYSP** statements prepared for an older version of the queue manager. It might also occur using a **CSQ6SYSP** parameter module built with an older version of the SCSQMACS macros.

#### **OTMACON**

OTMA parameters. This keyword takes five positional parameters::

**OTMACON** = ( **Group**,**Member**,**Druexit**,**Age**,**Tpipepfx** )

**Group** This is the name of the XCF group to which this particular instance of IBM MQ belongs.

It can be 1 through 8 characters long and must be entered in uppercase characters.

The default is blanks, which indicates that IBM MQ must not attempt to join an XCF group.

#### **Member**

This is the member name of this particular instance of IBM MQ within the XCF group.

It can be 1 through 16 characters long and must be entered in uppercase characters.

The default is the 4-character queue manager name.

#### **Druexit**

This specifies the name of the OTMA destination resolution user exit to be run by IMS.

It can be 1 through 8 characters long.

The default is DFSYDRU0.

This parameter is optional; it is required if IBM MQ is to receive messages from an IMS application that was not started by IBM MQ. The name must correspond to the destination resolution user exit coded in the IMS system. For more information see "Using OTMA exits in IMS" on page 1305.

**Age** This represents the length of time, in seconds, that a user ID from IBM MQ is considered previously verified by IMS.

It can be in the range zero through 2 147 483 647.

The default is 2 147 483 647.

You are recommended to set this parameter in conjunction with the `interval` parameter of the `ALTER SECURITY` command to maintain consistency of security cache settings across the mainframe.

### **Tpipepfx**

This represents the prefix to be used for Tpipe names.

It comprises three characters; the first character is in the range A through Z, subsequent characters are A through Z or 0 through 9. The default is CSQ.

This is used each time IBM MQ creates a Tpipe; the rest of the name is assigned by IBM MQ. You cannot set the full Tpipe name for any Tpipe created by IBM MQ.

### **QINDEXBLD**

Determines whether queue manager restart waits until all queue indexes are rebuilt, or completes before all indexes are rebuilt.

**WAIT** Queue manager restart waits for all queue index builds to be completed. This means that no applications are delayed during normal IBM MQ API processing while the index is created, as all indexes are created before any applications can connect to the queue manager.

This is the default.

### **NOWAIT**

The queue manager can restart before all queue index building is completed.

### **QMCCSID**

Specifies the default coded character set identifier that the queue manager (and therefore distributed queuing) is to use.

Specify a value in the range zero through 65535. The value must represent an EBCDIC codepage listed as a native z/OS codepage for your chosen language in National languages.

Zero, which is the default value, means use the CCSID currently set or, if none is set, use CCSID 500. This means that if you have explicitly set the CCSID to any non-zero value, you cannot reset it by setting QMCCSID to zero; you must now use the correct non-zero CCSID. If QMCCSID is zero, you can check what CCSID is actually in use by issuing the command `DISPLAY QMGR CCSID`.

On distributed platforms, use the `ALTER QMGR` command.

### **QSGDATA**

Queue-sharing group data. This keyword takes five positional parameters:

**QSGDATA=( Qsgname, Dsgname, Db2name, Db2serv , Db2b1ob)**

#### **Qsgname**

This is the name of the queue-sharing group to which the queue manager belongs.

See Rules for naming IBM MQ objects for valid characters. The name:

- Can be 1 through 4 characters long
- Must not start with a numeric
- Must not end in @.

This is because, for implementation reasons, names of less than four characters are padded internally with @ symbols,

The default is blanks, which indicates that the queue manager is not a member of any queue-sharing group.

**Dsgname**

This is the name of the Db2 data-sharing group to which the queue manager is to connect.

It can be 1 through 8 characters long and must be entered in uppercase characters.

The default is blanks, which indicates that you are not using queue-sharing groups.

**Db2name**

This is the name of the Db2 subsystem or group attachment to which the queue manager is to connect.

It can be 1 through 4 characters long and must be entered in uppercase characters.

The default is blanks, which indicates that you are not using queue-sharing groups.

**Note:** The Db2 subsystem (or group attachment) must be in the Db2 data-sharing group specified in the **Dsgname**, and all queue managers must specify the same Db2 data-sharing group.

**Db2serv**

This is the number of server tasks used for accessing Db2.

It can be in the range 4 through 10.

The default is 4.

**Db2blob**

This is the number of Db2 tasks used for accessing Binary Large Objects (BLOBs).

It can be in the range 4 through 10.

The default is 4.

If you specify one of the name parameters (that is, **Qsgname**, **Dsgname**, or **Db2name** ), you must enter values for the other names, otherwise IBM MQ fails.

**RESAUDIT**

Specifies whether RACF audit records are written for RESLEVEL security checks performed during connection processing.

Specify one of:

**NO** RESLEVEL auditing is not performed.

**YES** RESLEVEL auditing is performed.

The default is YES.

**ROUTCDE**

Specifies the default z/OS message routing code assigned to messages that are not sent in direct response to an MQSC command.

Specify one of:

1. A value in the range 1 through 16, inclusive.
2. A list of values, separated by a comma and enclosed in parentheses. Each value must be in the range 1 through 16, inclusive.

The default is 1.

For more information about z/OS routing codes, see the *MVS Routing and Descriptor Codes* manual.

**SERVICE**

This field is reserved for use by IBM.

**SMFACCT**

Specifies whether IBM MQ sends accounting data to SMF automatically when the queue manager starts.

Specify one of:

**NO** Do not start gathering accounting data automatically.

**YES** Start gathering accounting data automatically for the default class 1.

**integers**

A list of classes for which accounting is gathered automatically in the range 1 through 4.

The default is NO.

**SMFSTAT**

Specifies whether to gather SMF statistics automatically when the queue manager starts.

Specify one of:

**NO** Do not start gathering statistics automatically.

**YES** Start gathering statistics automatically for the default class 1.

**integers**

A list of classes for which statistics are gathered automatically in the range 1 through 4.

The default is NO.

**SPLCAP**

The security policy capability enables higher level of message security through policies that control whether messages are signed, or encrypted, as they are written and read from queues.

Its use is licensed by a separately installed product, IBM MQ Advanced Message Security (AMS), which supplies an enabling module in the SDRQAUTH library.

Security policy processing is enabled for this queue manager, by configuring SPLCAP with one of the following values:

**NO** The capability to implement message security policies for queues is not enabled during queue manager initialization.

**YES** Message security capabilities are enabled during queue manager initialization.

If this control is set, the queue manager attempts to load the license enabling module from SDRQAUTH during initialization, and start an additional address space (AMSM).

The queue manager does not start unless AMS is licensed, and the necessary configuration for message security is in place.

The default is NO.

**STATIME**

Specifies the default time, in minutes, between consecutive gatherings of statistics.

Specify a number in the range zero through 1440.

If you specify a value of zero, both statistics data and accounting data is collected at the SMF data collection broadcast. See Using System Management Facility for information about setting this.

The default is 30.

**TRACSTR**

Specifies whether global tracing is to start automatically.

Specify one of:

**NO** Do not start global tracing automatically.

**YES** Start global tracing automatically for the default class, class 1.

**integers**

A list of classes for which global tracing is to be started automatically in the range 1 through 4.

\* Start global trace automatically for all classes.

The default is NO if you do not specify the keyword in the macro.

**Note:** The supplied default system parameter load module (CSQZPARM) has TRACSTR=YES (set in the assembler module CSQFSYSP). If you do not want to start tracing automatically, either create your own system parameter module, or issue the STOP TRACE command after the queue manager has started.

For details about the STOP TRACE command, see STOP TRACE.

**TRACTBL**

Specifies the default size, in 4 KB blocks, of trace table where the global trace facility stores IBM MQ trace records.

Specify a value in the range 1 through 999.

The default is 99. This is equivalent to 396 KB.

**Note:** Storage for the trace table is allocated in the ECSA. Therefore, you must select this value with care.

**WLMTIME**

Specifies the time (in minutes or seconds depending on the value of WLMTIMU) between each scan of the indexes for WLM-managed queues.

Specify a value in the range 1 through 9999.

The default is 30.

**WLMTIMU**

Time units used with the WLMTIME parameter.

Specify one of :

**MINS** WLMTIME represents a number of minutes.

**SECS** WLMTIME represents a number of seconds.

The default is MINS.

**Related reference:**

“Using CSQ6LOGP”

Use this topic as a reference for how to specify logging options using CSQ6LOGP.

“Using CSQ6ARVP” on page 1235

Use this topic as a reference for how to specify your archiving environment using CSQ6ARVP

**Using CSQ6LOGP:**

Use this topic as a reference for how to specify logging options using CSQ6LOGP.

Use CSQ6LOGP to establish your logging options.

The default parameters for CSQ6LOGP, and whether you can alter each parameter using the `../com.ibm.mq.ref.adm.doc/q086640_dita` command, are shown in Default values of CSQ6LOGP parameters. If you need to change any of these values, refer to the detailed descriptions of the parameters.

*Table 153. Default values of CSQ6LOGP parameters*

Parameter	Description	Default value	SET command
COMPLOG	Controls whether log compression is enabled.	NONE	X
DEALLCT	Length of time an archive tape unit remains unused before it is deallocated.	zero	X
INBUFF	Size of input buffer storage for active and archive log data sets.	60 KB	-
MAXARCH	Maximum number of archive log volumes that can be recorded.	500	X
MAXCNOFF	Maximum number of CSQJOFF7 offload tasks that can be run in parallel.	31	-
MAXRTU	Maximum number of dedicated tape units allocated to read archive log tape volumes concurrently.	2	X
OFFLOAD	Archiving on or off.	YES (ON)	-
OUTBUFF	Size of output buffer storage for active and archive log data sets.	4 000 KB	-
TWOACTV	Single or dual active logging.	YES (dual)	-
TWOARCH	Single or dual archive logging.	YES (dual)	-
TWOBSDS	Single or dual BSDS.	YES (dual BSDS)	-
WRTHRSH	Number of output buffers to be filled before they are written to the active log data sets.	20	X

**COMPLOG**

Specifies whether log compression is enabled.

Specify either:

**NONE**

Log compression is not enabled.

**RLE** Log compression is enabled using run-length encoding.

**ANY** The queue manager selects the compression algorithm that gives the greatest degree of log record compression. This option results in RLE compression.

The default is NONE.

For more details about log compression, see Log compression.



## DEALLCT

Specifies the length of time, in minutes, that an archive read tape unit is allowed to remain unused before it is deallocated.

Specify one of the following:

- Time, in minutes, in the range zero through 1440
- NOLIMIT

Specifying 1440 or NOLIMIT means that the tape unit is never deallocated.

The default is zero.

When archive log data is being read from tape, it is recommended that you set this value high enough to allow IBM MQ to optimize tape handling for multiple read applications.

## INBUFF

Specifies the size, in kilobytes, of the input buffer for reading the active and archive logs during recovery. Use a decimal number in the range 28 through 60. The value specified is rounded up to a multiple of 4.

The default is 60 KB.

Suggested settings:

**Test system** 28 KB

**Production system** 60 KB

Set this to the maximum for best log read performance.

## MAXARCH

Specifies the maximum number of archive log volumes that can be recorded in the BSDS. When this number is exceeded, recording begins again at the start of the BSDS.

Use a decimal number in the range 10 through 1000.

The default is 500.

Suggested settings:

**Test system** 500 (default)

**Production system** 1 000

Set this to the maximum so that the BSDS can record as many logs as possible.

For information about the logs and BSDS, see *Managing IBM MQ resources*.

## MAXCNOFF

Specifies the number of CSQJOFF7 offload tasks that can be run in parallel.

This allows a queue manager, or queue managers, to be tuned such that they will not use all the available tape units.

Instead the queue manager waits until a CSQJOFF7 offload task has completed before trying to allocate any new archive data sets.

If the queue manager is archiving to tape, set this parameter so that the number of concurrent tape requests should not equal, or exceed, the number of tape units available, otherwise the system might hang.

Note that if dual archiving is in use, then each offload task performs both archives, so the parameter needs to be set accordingly. For example if the queue manager is dual archiving to tape, a value of MAXCNOFF=2 would allow up to two active logs to be archived concurrently to four tapes.

If several queue managers are sharing the tape units, you should set the MAXCNOFF for each queue manager accordingly.

The default value is 31.

Specify a value in the range 1 through 31.

#### **MAXRTU**

Specifies the maximum number of dedicated tape units that can be allocated to read archive log tape volumes concurrently.

This parameter and the DEALLCT parameter allow IBM MQ to optimize archive log reading from tape devices.

Specify a value in the range 1 through 99.

The default is 2.

It is recommended that you set the value to be at least one less than the number of tape units available to IBM MQ. If you do otherwise, the offload process could be delayed, which could affect the performance of your system. For maximum throughput during archive log processing, specify the largest value possible for this option, remembering that you need at least one tape unit for offload processing.

#### **OFFLOAD**

Specifies whether archiving is on or off.

Specify either:

**YES** Archiving is on

**NO** Archiving is off

The default is YES.

**Attention:** Do **not** switch archiving off unless you are working in a test environment. If you do switch it off, you cannot guarantee that data will be recovered in the event of a system or transaction failure.

#### **OUTBUFF**

Specifies the total size, in kilobytes, of the storage to be used by IBM MQ for output buffers for writing the active and archive log data sets. Each output buffer is 4 KB.

The parameter must be in the range 80 through 4000. The value specified is rounded up to a multiple of 4. Values between 40 and 80 will be accepted for compatibility reasons, and are treated as a value of 80.

The default is 4 000 KB.

Suggested settings:

<b>Test system</b>	400 KB
<b>Production system</b>	4 000 KB

Set this value to the maximum to avoid running out of log output buffers.

#### **TWOACTV**

Specifies single or dual active logging.

Specify either:

**NO** Single active logs

**YES** Dual active logs

The default is YES.

For more information about the use of single and dual logging, see Managing IBM MQ resources.

#### **TWOARCH**

Specifies the number of archive logs that IBM MQ produces when the active log is offloaded.

Specify either:

**NO** Single archive logs

**YES** Dual archive logs

The default is YES.

Suggested settings:

<b>Test system</b>	NO
<b>Production system</b>	YES (default)

For more information about the use of single and dual logging, see Managing IBM MQ resources.

#### **TWOBSDS**

Specifies the number of bootstrap data sets.

Specify either:

**NO** Single BSDDS

**YES** Dual BSDDS

The default is YES.

For more information about the use of single and dual logging, see Managing IBM MQ resources.

#### **WRTHRSH**

Specifies the number of 4 KB output buffers to be filled before they are written to the active log data sets.

The larger the number of buffers, the less often the write takes place, and this improves the performance of IBM MQ. The buffers might be written before this number is reached if significant events, such as a commit point, occur.

Specify the number of buffers in the range 1 through 256.

The default is 20.

#### **Related reference:**

“Using CSQ6SYSP” on page 1223

Use this topic as a reference for how to set system parameters using CSQ6SYSP.

“Using CSQ6ARVP”

Use this topic as a reference for how to specify your archiving environment using CSQ6ARVP

#### **Using CSQ6ARVP:**

Use this topic as a reference for how to specify your archiving environment using CSQ6ARVP

Use CSQ6ARVP to establish your archiving environment.

The default parameters for CSQ6ARVP, and whether you can alter each parameter using the SET ARCHIVE command, are shown in Table 154 on page 1236. If you need to change any of these values, refer to the detailed descriptions of the parameters. For more information about planning your archive storage, see Planning on z/OS.

Table 154. Default values of CSQ6ARVP parameters

Parameter	Description	Default value	SET command
ALCUNIT	Units in which primary and secondary space allocations are made.	BLK (blocks)	X
ARCPFX1	Prefix for first archive log data set name.	CSQARC1	X
ARCPFX2	Prefix for second archive log data set name.	CSQARC2	X
ARCRETN	The retention period of the archive log data set in days.	9999	X
ARCWRTC	List of route codes for messages to the operator about archive log data sets.	1,3,4	X
ARCWTOR	Whether to send message to operator and wait for reply before trying to mount an archive log data set.	YES	X
BLKSIZE	Block size of archive log data set.	28 672	X
CATALOG	Whether archive log data sets are cataloged in the ICF.	NO	X
COMPACT	Whether archive log data sets should be compacted.	NO	X
PRIQTY	Primary space allocation for DASD data sets.	25 715	X
PROTECT	Whether archive log data sets are protected by ESM profiles when the data sets are created.	NO	X
QUIESCE	Maximum time, in seconds, allowed for quiesce when ARCHIVE LOG with MODE(QUIESCE) specified.	5	X
SECQTY	Secondary space allocation for DASD data sets. See the ALCUNIT parameter for the units to be used.	540	X
TSTAMP	Whether the archive data set name should include a time stamp.	NO	X
UNIT	Device type or unit name on which the first copy of archive log data sets is stored.	TAPE	X
UNIT2	Device type or unit name on which the second copy of archive log data sets is stored.	Blank	X

#### ALCUNIT

Specifies the unit in which primary and secondary space allocations are made.

Specify one of:

**CYL** Cylinders

**TRK** Tracks

**BLK** Blocks

You are recommended to use BLK because it is independent of the device type.

The default is BLK.

If free space on the archive DASD volumes is likely to be fragmented, you are recommended to specify a smaller primary extent and allow expansion into secondary extents. For more information about space allocation for active logs, refer to the Planning on z/OS.

#### ARCPFX1

Specifies the prefix for the first archive log data set name.

See the TSTAMP parameter for a description of how the data sets are named and for restrictions on the length of ARCPFX1.

This parameter cannot be left blank.

The default is CSQARC1.

You might need to authorize the userid associated with the IBM MQ queue manager address space to create archive logs with this prefix.

#### **ARCPFX2**

Specifies the prefix for the second archive log data set name.

See the TSTAMP parameter for a description of how the data sets are named and for restrictions on the length of ARCPFX2.

This parameter cannot be blank even if the TWOARCH parameter is specified as NO.

The default is CSQARC2.

You might need to authorize the userid associated with the IBM MQ queue manager address space to create archive logs with this prefix.

#### **ARCRETN**

Specifies the retention period, in days, to be used when the archive log data set is created.

The parameter must be in the range zero through 9999.

The default is 9999.

Suggested settings:

<b>Test system</b>	3
	In a test system, archive logs are probably not required over long periods.
<b>Production system</b>	9 999 (default)
	Set this value high to effectively switch automatic archive log deletion off.

For more information about discarding archive log data sets, see *Administering IBM MQ for z/OS*.

#### **ARCWRTC**

Specifies the list of z/OS routing codes for messages about the archive log data sets to the operator. This field is ignored if ARCWTOR is set to NO.

Specify up to 14 routing codes, each with a value in the range 1 through 16. You must specify at least one code. Separate codes in the list by commas, not by blanks.

The default is the list of values: 1,3,4.

For more information about z/OS routing codes, see the *MVS Routing and Descriptor Codes* manual.

#### **ARCWTOR**

Specifies whether a message is to be sent to the operator and a reply is received before attempting to mount an archive log data set.

Other IBM MQ users might be forced to wait until the data set is mounted, but they are not affected while IBM MQ is waiting for the reply to the message.

Specify either:

**YES** The device needs a long time to mount archive log data sets. For example, a tape drive.

**NO** The device does not have long delays. For example, DASD.

The default is YES.

Suggested settings:

**Test system** NO

**Production system** YES (default)

This is dependent on operational procedures. If tape robots are used, NO might be more appropriate.

### **BLKSIZE**

Specifies the block size of the archive log data set. The block size you specify must be compatible with the device type you specify in the UNIT parameter.

The parameter must be in the range 4 097 through 28 672. The value you specify is rounded up to a multiple of 4 096.

The default is 28 672.

This parameter is ignored for data sets that are managed by the storage management subsystem (SMS).

If the archive log data set is written to DASD, you are recommended to choose the maximum block size that allows 2 blocks for each track. For example, for a 3390 device, you should use a block size of 24 576.

If the archive log data set is written to tape, specifying the largest possible block size improves the speed of reading the archive log. You should use a block size of 28 672.

Suggested settings:

**Test system** Use the block size recommendation depending on the media used for archive logs.

That is, for disk 24 576, and tape 28 672.

**Production system** Use the block size recommendation depending on the media used for archive logs.

That is, for disk 24 576, and tape 28 672.

### **CATALOG**

Specifies whether archive log data sets are cataloged in the primary integrated catalog facility (ICF) catalog.

Specify either:

**NO** Archive log data sets are not cataloged

**YES** Archive log data sets are cataloged

The default is NO.

All archive log data sets allocated on DASD must be cataloged. If you archive to DASD with the CATALOG parameter set to NO, message CSQJ072E is displayed each time an archive log data set is allocated, and IBM MQ catalogs the data set.

Suggested settings:

**Test system** YES

**Production system** YES, when archives are allocated on DASD

### **COMPACT**

Specifies whether data written to archive logs is to be compacted. This option applies only to a 3480 or 3490 device that has the improved data recording capability (IDRC) feature. When this feature is turned on, hardware in the tape control unit writes data at a much higher density than normal, allowing for more data on each volume. Specify NO if you do not use a 3480 device with the IDRC feature or a 3490 base model, except for the 3490E. Specify YES if you want the data to be compacted.

Specify either:

**NO** Do not compact the data sets

**YES** Compact the data sets

The default is NO.

Specifying YES adversely affects performance. Also be aware that data compressed to tape can be read only using a device that supports the IDRC feature. This can be a concern if you have to send archive tapes to another site for remote recovery.

Suggested settings:

**Test system** Not applicable

**Production system** NO (default)

This applies to 3480 and 3490 IDR compression only. Setting this to YES might degrade archive log read performance during recovery and restart; however, it does not affect writing to tape.

### **PRIQTY**

Specifies the primary space allocation for DASD data sets in ALCUNITs.

The value must be greater than zero.

The default is 25 715.

This value must be sufficient for a copy of either the log data set or its corresponding BSDS, whichever is the larger. To determine the necessary value, follow this procedure:

1. Determine the number of active log records allocated ( c ) as explained in "Task 14: Create the bootstrap and log data sets" on page 1218.
2. Determine the number of 4096 byte blocks in each archive log block:

$$d = \text{BLKSIZE} / 4096$$

where BLKSIZE is the rounded up value.

3. If ALCUNIT=BLK:

$$\text{PRIQTY} = \text{INT}(c / d) + 1$$

where INT means round down to an integer.

If ALCUNIT=TRK:

$$\text{PRIQTY} = \text{INT}(c / (d * \text{INT}(e/\text{BLKSIZE}))) + 1$$

where e is the number of bytes for each track (56664 for a 3390 device) and INT means round down to an integer.

If ALCUNIT=CYL:

$$\text{PRIQTY} = \text{INT}(c / (d * \text{INT}(e/\text{BLKSIZE}) * f)) + 1$$

where  $f$  is the number of tracks for each cylinder (15 for a 3390 device) and INT means round down to an integer.

For information about how large to make your log and archive data sets, see “Task 14: Create the bootstrap and log data sets” on page 1218 and “Task 15: Define your page sets” on page 1219.

Suggested settings:

<b>Test system</b>	1 680
	Sufficient to hold the entire active log, that is: $10\ 080 / 6 = 1\ 680$ blocks
<b>Production system</b>	Not applicable when archiving to tape.

If free space on the archive DASD volumes is likely to be fragmented, you are recommended to specify a smaller primary extent and allow expansion into secondary extents. For more information about space allocation for active logs, refer to the Planning on z/OS.

#### PROTECT

Specifies whether archive log data sets are to be protected by discrete ESM (external security manager) profiles when the data sets are created.

Specify either:

**NO** Profiles are not created.

**YES** Discrete data set profiles are created when logs are offloaded. If you specify YES:

- ESM protection must be active for IBM MQ.
- The user ID associated with the IBM MQ queue manager address space must have authority to create these profiles.
- The TAPEVOL class must be active if you are archiving to tape.

Otherwise, offloading fails.

The default is NO.

#### QUIESCE

Specifies the maximum time in seconds allowed for the quiesce when an ARCHIVE LOG command is issued with MODE(QUIESCE) specified.

The parameter must be in the range 1 through 999.

The default is 5.

#### SECQTY

Specifies the secondary space allocation for DASD data sets in ALCUNITs. The secondary extent can be allocated up to 15 times; see the *z/OS MVS JCL Reference* and *z/OS MVS JCL User's Guide* for details.

The parameter must be greater than zero.

The default is 540.

#### TSTAMP

Specifies whether the archive log data set name has a time stamp in it.

Specify either:

**NO** Names do not include a time stamp. The archive log data sets are named:



*arcpfxi.A nnnnnnn*

Where *arcpfxi* is the data set name prefix specified by ARCPFX1 or ARCPFX2. *arcpfxi* can have up to 35 characters.

**YES** Names include a time stamp. The archive log data sets are named:

*arcpfxi.cyyddd.T hhmsst.A nnnnnnn*

where *c* is 'D' for the years up to and including 1999 or 'E' for the year 2000 and later, and *arcpfxi* is the data set name prefix specified by ARCPFX1 or ARCPFX2. *arcpfxi* can have up to 19 characters.

**EXT** Names include a time stamp. The archive log data sets are named:

*arcpfxi.D yyyyddd.T hhmsst.A nnnnnnn*

Where *arcpfxi* is the data set name prefix specified by ARCPFX1 or ARCPFX2. *arcpfxi* can have up to 17 characters.

The default is NO.

#### **UNIT**

Specifies the device type or unit name of the device that is used to store the first copy of the archive log data set.

Specify a device type or unit name of 1 through 8 alphanumeric characters. The first character must be alphabetic.

This parameter cannot be blank.

The default is TAPE.

If you archive to DASD, you can specify a generic device type with a limited volume range, for example, UNIT=3390.

If you archive to DASD, make sure that:

- The primary space allocation is large enough to contain all the data from the active log data sets.
- The archive log data set catalog option (CATALOG) is set to YES.
- You have used a proper value for BLKSIZE.

If you archive to TAPE, IBM MQ can extend to a maximum of 20 volumes.

Suggested settings:

<b>Test system</b>	DASD
<b>Production system</b>	TAPE

For more information about choosing a location for archive logs, refer to the Planning on z/OS.

#### **UNIT2**

Specifies the device type or unit name of the device that is used to store the second copy of the archive log data sets.

Specify a device type or unit name of 1 through 8 alphanumeric characters. The first character must be alphabetic. If this parameter is blank, the value set for the UNIT parameter is used.

The default is blank.

**Related reference:**

“Using CSQ6SYSP” on page 1223

Use this topic as a reference for how to set system parameters using CSQ6SYSP.

“Using CSQ6LOGP” on page 1232

Use this topic as a reference for how to specify logging options using CSQ6LOGP.

**Task 18: Tailor the channel initiator parameters**

Use ALTER QMGR to customize the channel initiator to suit your requirements.

- *Repeat this task for each IBM MQ queue manager, as required.*
- *You must perform this task when migrating from a previous version.*

A number of queue manager attributes control how distributed queuing operates. Set these attributes using the MQSC command ALTER QMGR. The initialization data set sample thlqual.SCSQPROC(CSQ4INYG) contains some settings that you can customize. For more information, see ALTER QMGR.

The values of these parameters are displayed as a series of messages each time you start the channel initiator.

**The relationship between adapters, dispatchers, and maximum number of channels**

The ALTER QMGR parameters CHIADAPS and CHIDISPS define the number of task control blocks (TCBs) used by the channel initiator. CHIADAPS (adapter) TCBs are used to make IBM MQ API calls to the queue manager. CHIDISPS (dispatcher) TCBs are used to make calls to the communications network.

The ALTER QMGR parameter MAXCHL influences the distribution of channels over the dispatcher TCBs.

**CHIDISPS**

If you have a small number of channels use the default value.

One task for each processor optimizes system performance. As dispatcher tasks are CPU intensive, the principle is to keep as few tasks as busy as possible, so that the time taken to find and start threads is minimized.

CHIDISPS(20) is suitable for systems with more than 100 channels. There is unlikely to be any significant disadvantage in having CHIDISPS(20) where this is more dispatcher TCBs than necessary.

As a guideline, if you have more than 1000 channels, allow one dispatcher for every 50 current channels. For example, specify CHIDISPS(40) to handle up to 2000 active channels.

If you are using TCP/IP, the maximum number of dispatchers used for TCP/IP channels is 100, even if you specify a larger value in CHIDISPS.

**CHIADAPS**

Each IBM MQ API call to the queue manager is independent of any other and can be made on any adapter TCB. Calls using persistent messages can take much longer than those for nonpersistent messages because of log I/O. Thus a channel initiator processing a large number of persistent messages across many channels may need more than the default 8 adapter TCBs for optimum performance. This is particularly so where achieved batchsize is small, because end of batch processing also requires log I/O, and where thin client channels are used.

The suggested value for a production environment is CHIADAPS(30). Using more than this is unlikely to give any significant extra benefit, and there is unlikely to be any significant disadvantage in having CHIADAPS(30) if this is more adapter TCBs than necessary.

**MAXCHL**

Each channel is associated with a particular dispatcher TCB at channel start and remains

associated with that TCB until the channel stops. Many channels can share each TCB. MAXCHL is used to spread channels across the available dispatcher TCBs. The first (  $\text{MIN}(\text{MAXCHL} / \text{CHIDISPS}), 10$  ) channels to start are associated with the first dispatcher TCB, and so on, until all dispatcher TCBs are in use.

The effect of this for small numbers of channels and a large MAXCHL is that channels are NOT evenly distributed across dispatchers. For example, if you set CHIDISPS(10) and left MAXCHL at its default value of 200 but had only 50 channels, five dispatchers would be associated with 10 channels each and five would be unused. We suggest setting MAXCHL to the number of channels actually to be used where this is a small fixed number.

If you change this queue manager property, you must also review the ACTCHL, LU62CHL, and TCPCHL queue manager properties to ensure that the values are compatible. See Queue manager parameters for a full description of these properties, and their relationship.

## Setting up your z/OS UNIX System Services environment for channel initiators

The channel initiator (CHINIT) uses OMVS threads. Review the OMVS configuration parameters before creating a new CHINIT, or modifying the number of dispatchers or SSLTASKS.

Each CHINIT uses  $3 + \text{CHIDISP} + \text{SSLTASKS}$  OMVS threads. These contribute to the total number of OMVS threads used in the LPAR, and towards the number of threads used by CHINIT started task user ID.

You can use the **D OMVS,L** and review the current usage, highwater usage, and system limit of MAXPROCSYS (the maximum number of processes that the system allows).

If you are adding a new CHINIT or increasing the values of CHIDISPS or SSLTASKS then you must calculate the increase in threads and review the impact on the MAXPROCSYS values. You can use the **SETOMVS** command to dynamically change the MAXPROCSYS, or update the BPXPRCxx parmlib value or both.

The OMVS parameter MAXPROCUSER is the number of OMVS threads a single OMVS user, that is with the same UID, can have. The threads count towards this value. So if you have 2 CHINITs with the same started task user ID, with 10 dispatchers and 3 SSLTASKS each then there are  $2 * (3 + 10 + 3) = 32$  threads for the OMVS uid.

You can display the default MAXPROCUSER by issuing the **D OMVS,0** command and you can use the **SETOMVS** command to dynamically change the MAXPROCUSER, or update the BPXPRCxx parmlib value or both.

You can override this value on a per user basis with the RACF command **ALTUSER userid OMVS(PROCUSERMAX(nnnn))** or equivalent.

To start the channel initiator, issue the following command:

```
START CHINIT
```

To ensure that the channel initiator has started successfully, check that there is no ICH408I error in the xxxxCHIN(ssidCHIN) job log.

## Related concepts:

“Task 19: Set up Batch, TSO, and RRS adapters”

Make the adapters available to applications by adding libraries to appropriate STEPLIB concatenations. To cater for SNAP dumps issued by an adapter, allocate a CSQSNAP DDname. Consider using CSQBDEFV to improve the portability of your application programs

## Related information:

Channel initiator statistics data records

## Task 19: Set up Batch, TSO, and RRS adapters

Make the adapters available to applications by adding libraries to appropriate STEPLIB concatenations. To cater for SNAP dumps issued by an adapter, allocate a CSQSNAP DDname. Consider using CSQBDEFV to improve the portability of your application programs

- Repeat this task for each IBM MQ queue manager as required.
- You might need to perform this task when migrating from a previous version.

To make the adapters available to batch and other applications using batch connections, add the following IBM MQ libraries to the STEPLIB concatenation for your batch application :

- thlqual.SCSQANL *x*
- thlqual.SCSQAUTH

where *x* is the language letter for your national language. (You do not need to do this if the libraries are in the LPA or the link list.)

For TSO applications add the libraries to the STEPLIB concatenation in the TSO logon procedure or activate them using the TSO command TSOLIB.

If the adapter detects an unexpected IBM MQ error, it issues an z/OS SNAP dump to DDname CSQSNAP, and issues reason code MQRD\_UNEXPECTED\_ERROR to the application. If the CSQSNAP DD statement is not in the application JCL or CSQSNAP is not allocated to a data set under TSO, no dump is taken. If this happens, you could include the CSQSNAP DD statement in the application JCL or allocate CSQSNAP to a data set under TSO and rerun the application. However, because some problems are intermittent, it is recommended that you include a CSQSNAP statement in the application JCL or allocate CSQSNAP to a data set in the TSO logon procedure to capture the reason for failure at the time it occurs.

The supplied program CSQBDEFV improves the portability of your application programs. In CSQBDEFV, you can specify the name of a queue manager, or queue sharing group, to be connected to rather than specifying it in the MQCONN or MQCONNX call in an application program. You can create a new version of CSQBDEFV for each queue manager, or queue sharing group. To do this, follow these steps:

1. Copy the IBM MQ assembler program CSQBDEFV from thlqual.SCSQASMS to a user library.
2. The supplied program contains the default subsystem name CSQ1. You can retain this name for testing and installation verification. For production subsystems, you can change the NAME=CSQ1 to your one- to four-character subsystem name, or use CSQ1.

If you are using queue-sharing groups, you can specify a queue-sharing group name instead of CSQ1. If you do this, the program issues a connect request to an active queue manager within that group.

3. Assemble and link-edit the program to produce the CSQBDEFV load module. For the assembly, include the library thlqual.SCSQMACS in your SYSLIB concatenation; use the link-edit parameters RENT,AMODE=31,RMODE=ANY. This is shown in the sample JCL in thlqual.SCSQPROC(CSQ4DEFV). Then include the load library in the z/OS Batch or the TSO STEPLIB, ahead of thlqual.SCSQAUTH.

## Related concepts:

“Task 20: Set up the operations and control panels”

To set up the operations and control panels you must first set up the libraries that contain the required panels, EXECs, messages, and tables. To do this, you must take into account which national language feature is to be used for the panels. When you have done this, you can optionally update the main ISPF menu for IBM MQ operations and control panels and change the function key settings.

## Task 20: Set up the operations and control panels

To set up the operations and control panels you must first set up the libraries that contain the required panels, EXECs, messages, and tables. To do this, you must take into account which national language feature is to be used for the panels. When you have done this, you can optionally update the main ISPF menu for IBM MQ operations and control panels and change the function key settings.

- *You need to perform this task once for each z/OS system where you want to run IBM MQ.*
- *You might need to perform this task when migrating from a previous version.*

## Setting up the libraries

Follow these steps to set up the IBM MQ operations and control panels:

1. Ensure that all the libraries contained in your concatenations are either in the same format (F, FB, V, VB) and have the same block size, or are in order of decreasing block sizes. Otherwise, you might have problems trying to use these panels.
2. Include the library thlqual.SCSQEXEC in your SYSEXEC or SYSPROC concatenation or activate it using the TSO ALTLIB command. This library, which is allocated with a fixed-block 80 record format during installation, contains the required EXECs.

It is preferable to put the library into your SYSEXEC concatenation. However, if you want to put it in SYSPROC, the library must have a record length of 80 bytes.

3. Add thlqual.SCSQAUTH and thlqual.SCSQANLx to the TSO logon procedure STEPLIB or activate it using the TSO TSOLIB command, if it is not in the link list or the LPA.
4. You can either add the IBM MQ panel libraries permanently to your ISPF library setup, or allow them to be set up dynamically when the panels are used. For the former choice, you need to do the following:
  - a. Include the library containing the operations and control panel definitions in your ISPPLIB concatenation. The name is thlqual.SCSQPNLx, where x is the language letter for your national language.
  - b. Include the library containing the required tables in your ISPTLIB concatenation. The name is thlqual.SCSQTBLx, where x is the language letter for your national language.
  - c. Include the library containing the required messages in your ISPMLIB concatenation. The name is thlqual.SCSQMSGx, where x is the language letter for your national language.
  - d. Include the library containing the required load modules in your ISPLLIB concatenation. The name of this library is thlqual.SCSQAUTH.
5. Test that you can access the IBM MQ panels from the TSO Command Processor panel. This is usually option 6 on the ISPF/PDF Primary Options Menu. The name of the EXEC that you run is CSQOREXX. There are no parameters to specify if you have put the IBM MQ libraries permanently in your ISPF setup as in step 4. If you have not, use the following:

```
CSQOREXX thlqual langletter
```

where langletter is a letter identifying the national language to be used:

- C Simplified Chinese
- E U.S. English (mixed case)
- F French
- K Japanese
- U U.S. English (uppercase)

## Updating the ISPF menu

You can update the ISPF main menu to allow access to the IBM MQ operations and control panels from ISPF. The required setting for &ZSEL is:

```
CMD(%CSQOREXX thlqual langletter)
```

For information about thlqual and langletter, see Step 5 on page 1245.

For more details, see the *z/OS: ISPF Dialog Developer's Guide and Reference* manual.

## Updating the function keys and command settings

You can use the normal ISPF procedures for changing the function keys and command settings used by the panels. The application identifier is CSQO.

However, this is *not* recommended because the help information is not updated to reflect any changes that you have made.

### Related concepts:

“Task 21: Include the IBM MQ dump formatting member”

To be able to format IBM MQ dumps using the Interactive Problem Control System (IPCS), you must update some system libraries.

### Task 21: Include the IBM MQ dump formatting member

To be able to format IBM MQ dumps using the Interactive Problem Control System (IPCS), you must update some system libraries.

- *You need to perform this task once for each z/OS system where you want to run IBM MQ.*
- *You need to perform this task when migrating from a previous version.*

To be able to format IBM MQ dumps using the Interactive Problem Control System (IPCS), copy the data set thlqual.SCSQPROC(CSQ7IPCS) to SYS1.PARMLIB. You should not need to edit this data set.

If you have customized the TSO procedure for IPCS, thlqual.SCSQPROC(CSQ7IPCS) can be copied into any library in the IPCSPARM definition. See the *MVS IPCS Customization* manual for details on IPCSPARM.

You must also include the library thlqual.SCSQPNLA in your ISPPLIB concatenation.

To make the dump formatting programs available to your TSO session or IPCS job, you must also include the library thlqual.SCSQAUTH in your STEPLIB concatenation or activate it using the TSO TSOLIB command (even if it is already in the link list or LPA).

**Related concepts:**

“Task 22: Suppress information messages”

Your IBM MQ system might produce a large number of information messages. You can prevent selected messages being sent to the console or to the hardcopy log.

**Task 22: Suppress information messages**

Your IBM MQ system might produce a large number of information messages. You can prevent selected messages being sent to the console or to the hardcopy log.

- *You need to perform this task once for each z/OS system where you want to run IBM MQ.*
- *You do not need to perform this task when migrating from a previous version.*

If your IBM MQ system is heavily used, with many channels stopping and starting, a large number of information messages are sent to the z/OS console and hardcopy log. The IBM MQ - IMS bridge and buffer manager might also produce a large number of information messages.

If required, you can suppress some of these console messages by using the z/OS message processing facility list, specified by the MPFLSTxx members of SYS1.PARMLIB. The messages you specify still appear on the hardcopy log, but not on the console.

Sample thlqual.SCSQPROC(CSQ4MPFL) shows suggested settings for MPFLSTxx. See the *MVS Initialization and Tuning Reference* manual for more information about MPFLSTxx.

If you want to suppress selected information messages on the hardcopy log, you can use the z/OS installation exit IEAVMXIT. You can set the following bit switches ON for the required messages:

**CTXTRDTM**

Delete the message.

The message is not displayed on consoles or logged in hardcopy.

**CTXTESJL**

Suppress from job log.

The message does not go into the JES job log.

**CTXTNWTP**

Do not carry out WTP processing.

The message is not sent to a TSO terminal or to the system message data set of a batch job.

**Note:**

1. For full details on the other parameters, refer to the MVS Installation Exits documentation.
2. You are not recommended to suppress messages other than those in the suggested suppression list, CSQ4MPFL.

In addition you can specify the extra parameter:

**EXCLMSG**

Specifies a list of messages to be excluded from any log.

Messages in this list are not sent to the z/OS console and hardcopy log. See EXCLMSG in “Using CSQ6SYSP” on page 1223 for further information.

### Related concepts:

“Testing your queue manager on z/OS” on page 1252

When you have customized or migrated your queue manager, you can test it by running some of the sample applications shipped with IBM MQ.

### Task 23: Create procedures for Advanced Message Security

Each IBM MQ subsystem that is to be configured to use Advanced Message Security requires a cataloged procedure to start the AMS address space. You can create your own or use the IBM-supplied procedure library.

For each IBM MQ subsystem that is to be configured to use Advanced Message Security tailor a copy of sample procedure CSQ4AMSM. To do this, perform the following steps:

1. Copy the sample started task procedure *thlqual.SCSQPROC(CSQ4AMSM)* to your SYS1.PROCLIB or, if you are not using SYS1.PROCLIB, your procedure library. Name the procedure *xxxxAMSM*, where *xxxx* is the name of your IBM MQ subsystem. For example, *CSQ1AMSM* would be the AMS started task procedure for queue manager CSQ1.
2. Make a copy for each IBM MQ subsystem that you are going to use.
3. Tailor the procedures to your requirements using the instructions in the sample procedure CSQ4AMSM. You can also use symbolic parameters in the JCL to allow the procedure to be modified when it is started.
4. Review and optionally change the parameters passed to the AMS task using the Language Environment<sup>®</sup> *\_CEE\_ENVFILE* file. The sample *thlqual.SCSQPROC(CSQ40ENV)* lists the supported parameters.

**Note:** This task should be repeated for each IBM MQ queue manager.

### Task 24: Set up the started task user Advanced Message Security

The IBM MQ Advanced Message Security task requires a user ID that allows it to be known as a UNIX System Services process.

In addition, the users that the task works on behalf of must also have an appropriate definition of a UNIX UID (user ID) and GID (group ID) so these users are known as UNIX System Services users. For more information on defining UNIX System Services UIDs and GIDs, see *z/OS: Security Server RACF Security Administrator's Guide*.

*z/OS: UNIX System Services Planning* compares traditional UNIX security to z/OS security. The primary difference between traditional UNIX security and z/OS security is that the Kernel services support two levels of appropriate privileges: UNIX level and z/OS UNIX level.

Depending on your installation's security policy, the IBM MQ Advanced Message Security task can either run with superuser authority (*uid(0)*), or with its RACF identity permitted to the RACF FACILITY class *BPX.DAEMON* and *BPX.SERVER* profiles, as this task must be able to assume the RACF identity of its users.

If the latter method is used, or you have already activated the *BPX.DAEMON* or *BPX.SERVER* profiles, the IBM MQ Advanced Message Security task program (*thlqual.SCSQAUTH(CSQ0DSRV)*) must be located in RACF program-controlled libraries.

Review *z/OS: UNIX System Services Planning* to ensure that you understand the security differences between traditional UNIX security and z/OS UNIX security. This allows you to administer the IBM MQ Advanced Message Security task according to your installation's security policy for deploying and running privileged UNIX System Services processes.

For reference, the publications useful to this review are:

- *z/OS: UNIX System Services Planning*.



- *z/OS: Security Server RACF Security Administrator's Guide.*

**Note:** Choose the user ID for this task carefully because the IBM MQ Advanced Message Security recipient certificates are loaded into a key ring associated with this user ID. This consideration is discussed in Using certificates on z/OS .

The steps shown here describe how to set up the IBM MQ Advanced Message Security started task user. The steps use RACF commands as examples. If you are using a different security manager, you should use equivalent commands.

**Note:** The examples in this section assume that you have activated generic profile command processing for the RACF STARTED, FACILITY, and SURROGAT classes and generic profile checking. For more information on how RACF handles generic profiles, see *z/OS: Security Server RACF Command Language Reference*.

1. First define RACF user profiles for the IBM MQ Advanced Message Security started task user. These can be the same user.

```
ADDUSER WMQMSM NAME('IBM MQ Advanced Message Security user') OMVS (UID(0)) DFLTGRP(group)
```

Select a default 'group' as appropriate to your installation standards.

**Note:** If you do not want to grant USS superuser authority (UID(0)), then you must permit the IBM MQ Advanced Message Security user ID to the BPX.DAEMON and BPX.SERVER facility class profiles:

```
PERMIT BPX.DAEMON CLASS(FACILITY) ID(WMQMSM) ACCESS(READ)
```

and the IBM MQ Advanced Message Security task program ( *thlqual.SCSQAUTH(CSQ0DSRV)*) must be located in a RACF program-controlled library.

To make your SCSQAUTH library program controlled, you can use the following command:

```
RALTER PROGRAM * ADDMEM('thlqual.SCSQAUTH'//NOPADCHK) -or-
RALTER PROGRAM ** ADDMEM('thlqual.SCSQAUTH'//NOPADCHK)
SETROPTS WHEN(PROGRAM) REFRESH
```

You must also enable program control for the national language library ( *thlqual.SCSQANLx*) that is used by the IBM MQ Advanced Message Security task.

2. Determine if the RACF STARTED class is active. If it is not, activate the RACF STARTED class:
 

```
SETROPTS CLASSACT(STARTED)
```

3. Define a started class profile for the IBM MQ Advanced Message Security tasks, specifying the user IDs you selected or created in step 1:

```
RDEFINE STARTED qmgr AMSM.* STDATA(USER(WMQMSM))
```

where *qmgr* is the name of prefix of the started task name. For example, the started tasks may be named CSQ1AMSM. In this case, you would substitute *qmgr* AMSM.\* with CSQ1AMSM.\*.

The started task names must be named *qmgr* AMSM.\*.

4. Use the SETROPTS RACF command to refresh the in-storage RACLISTed started class profiles:
 

```
SETROPTS RACLIST(STARTED) REFRESH
```

5. The IBM MQ Advanced Message Security task temporarily assumes the identity of the host user ID of the client requestor during protection processing of IBM MQ messages. Therefore, it is necessary to define profiles in the SURROGAT class for each user ID that can make requests.

This can be done with a single generic profile if the RACF SURROGAT class is active. The check is ignored if the SURROGAT class is not active. The SURROGAT profiles needed are described in *z/OS: UNIX System Services Planning*.

To define profiles in the SURROGAT class:

- a. Activate the RACF SURROGAT class using the RACF SETROPTS command:

```
SETRPTS CLASSACT(SURROGAT)
```

- b. Activate generic profile processing for the RACF SURROGAT class:

```
SETRPTS GENERIC(SURROGAT)
```

- c. Activate generic profile command processing for the RACF SURROGAT class:

```
SETRPTS GENCMD(SURROGAT)
```

- d. Define a surrogate class generic profile:

```
RDEFINE SURROGAT BPX.SRV.* UACC(NONE)
```

- e. Permit the IBM MQ Advanced Message Security user ID to the generic SURROGAT class profile:

```
PERMIT BPX.SRV.* CLASS(SURROGAT) ID(WMQASM) ACCESS(UPDATE)
```

**Note:** You can define more specific profiles if you want to restrict specific users to be processed by the IBM MQ Advanced Message Security task, as described in *z/OS: UNIX System Services Planning*.

- f. Permit the IBM MQ Advanced Message Security user ID to the BPX.SERVER facility (if not already done in Creating the certificates and key rings ):

```
PERMIT BPX.SERVER CLASS(FACILITY) ID(WMQASM) ACCESS(READ)
```

6. The IBM MQ Advanced Message Security task uses the facilities provided by z/OS System SSL services to open SAF-managed key rings. The underlying System Authorization Facility (SAF) that accesses the contents of the key rings is controlled by RACF, or an equivalent security manager.

This service is the IRRSDL00 (R\_datalib) callable service. This callable service is protected with the same profiles used to protect the RACF RACDCERT commands that are defined to the RACF FACILITY class. Thus, the IBM MQ Advanced Message Security user ID must be permitted to the profiles using these commands:

- a. If you have not already done so, define a RACF generic profile to the RACF FACILITY class that protects the RACDCERT command and the IRRSDL00 callable service:

```
RDEFINE FACILITY IRR.DIGTCERT.* UACC(NONE)
SETRPTS RACLIST(FACILITY) REFRESH
```

- b. Grant authority to the started task user ID to the RACF generic profile:

```
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(WMQASM) ACC(READ)
```

Alternatively, you can grant READ access to the data service task user's keyring in the RDATA LIB class as follows:

```
PERMIT WMQASMD.DRQ.AMS.KEYRING.LST CLASS(RDATA LIB) ID(WMQASM) ACC(READ)
```

## Resource security for AMS

The started task user requires read authority to the SYSTEM.PROTECTION.POLICY.QUEUE.

The started task user requires authority to connect to the queue manager as a BATCH application. For further information, see Connection security profiles for batch connections.

## Task 25: Grant RACDCERT permissions to the security administrator for Advanced Message Security

Your IBM MQ Advanced Message Security security administrator requires authority to use the RACDCERT command to create and manage digital certificates.

Identify the appropriate user ID for this role and grant permission to use the RACDCERT command. For example:

```
PERMIT IRR.DIGTCERT.* CLASS(FACILITY) ID(admin) ACCESS(CONTROL)
SETOPTS RACLIST(FACILITY) REFRESH
```

where `admin` is the user ID of your IBM MQ Advanced Message Security security administrator.

## Task 26: Grant users resource permissions for IBM MQ Advanced Message Security

IBM MQ Advanced Message Security users require relevant resource permissions.

IBM MQ Advanced Message Security users, that is users that are putting or getting IBM MQ Advanced Message Security protected messages, require:

- An OMVS segment associated with their user id
- Permissions for IRR.DIGTCERT.LISTRING or RDATA LIB
- Permissions for ICSF class CSFSERV and CSFKEYS profiles

The IBM MQ Advanced Message Security task temporarily assumes the identity of its clients; that is, the task acts as a surrogate of the z/OS user ID of users of IBM MQ Advanced Message Security during the processing of IBM MQ messages to queues that are protected by IBM MQ Advanced Message Security.

In order for the task to assume the z/OS identity of a user, the client z/OS user ID must have a defined OMVS segment associated with its user profile.

As an administration aid, RACF provides the ability to define a default OMVS segment that may be associated with RACF user and group profiles. This default is used if the z/OS user ID or group profile does not have an OMVS segment explicitly defined. If you plan to have a large number of users using IBM MQ Advanced Message Security, you may choose to use this default rather than explicitly defining the OMVS segment for each user.

The *z/OS: Security Server RACF Security Administrator's Guide* contains the detailed procedure for defining default OMVS segments. Review the procedure as outlined in this publication to determine if the definition of default OMVS segments in RACF User and Group profiles is appropriate to your installation.

To grant READ permission to the IRR.DIGTCERT.LISTRING class facility to all IBM MQ Advanced Message Security users, issue this command:

```
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(READ)
```

or grant READ permission on a per user basis by issuing this command:

```
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(userid) ACCESS(READ)
```

where `userid` is the name of the IBM MQ Advanced Message Security user.

Alternatively, you can use the RDATA LIB class to grant access to specific keyrings (the RDATA LIB permissions take precedence over IRR.DIGTCERT.LISTRING permissions). For example:

```
PERMIT user.DRQ.AMS.KEYRING.LST CLASS(RDATA LIB) ID(user) ACC(READ)
```

If you are using ICSF-managed certificates and private keys, IBM MQ Advanced Message Security users require access to certain class CSFSERV and CSFKEYS profiles. This access is detailed in the following

table:

Table 155. Required user access to class CSFSERV and CSFKEYS profiles

Class	Profile	Permission
CSFSERV	CSFDSG	READ
CSFSERV	CSFPKE	READ
CSFSERV	CSFPKD	READ
CSFSERV	CSFDSV	READ
CSFKEYS	ICSF PKDS Label	READ

## Testing your queue manager on z/OS

When you have customized or migrated your queue manager, you can test it by running some of the sample applications shipped with IBM MQ.

You can compile and link-edit whichever of the other samples are appropriate to your installation using the sample JCL supplied.

See the following links for instructions on how to test your queue manager on z/OS:

- “Running the basic installation verification program”
- “Testing for queue-sharing groups” on page 1256
- “Testing for distributed queuing” on page 1257
- “Testing for C, C++, COBOL, PL/I, and CICS” on page 1261

### Related concepts:

“Configuring queue managers on z/OS” on page 1194

Use these instructions to configure queue managers on IBM MQ for z/OS.

### Related information:

IBM MQ for z/OS concepts

Planning your IBM MQ environment on z/OS

Administering IBM MQ for z/OS

## Running the basic installation verification program

After you have installed and customized IBM MQ, you can use the supplied installation verification program, CSQ4IVP1, to confirm that IBM MQ is operational.

The basic installation verification program is a batch assembler IVP that verifies the base IBM MQ without using the C, COBOL, or CICS samples.

The Batch Assembler IVP is link-edited by SMP/E and the load modules are shipped in library thlqual.SCSQLOAD.

After you have completed both the SMP/E APPLY step and the customization steps, run the Batch Assembler IVP.

See these sections for further details:

- Overview of the CSQ4IVP1 application
- Preparing to run CSQ4IVP1
- Running CSQ4IVP1
- Checking the results of CSQ4IVP1

## Overview of the CSQ4IVP1 application

CSQ4IVP1 is a batch application that connects to your IBM MQ subsystem and performs these basic functions:

- Issues IBM MQ calls
- Communicates with the command server
- Verifies that triggering is active
- Generates and deletes a dynamic queue
- Verifies message expiry processing
- Verifies message commit processing

## Preparing to run CSQ4IVP1

Before you run CSQ4IVP1:

1. Check that the IVP entries are in the CSQINP2 data set concatenation in the queue manager startup program. The IVP entries are supplied in member thlqual.SCSQPROC(CSQ4IVPQ). If not, add the definitions supplied in thlqual.SCSQPROC(CSQ4IVPQ) to your CSQINP2 concatenation. If the queue manager is currently running, you need to restart it so that these definitions can take effect.
2. The sample JCL, CSQ4IVPR, required to run the installation verification program is in library thlqual.SCSQPROC.

Customize the CSQ4IVPR JCL with the high-level qualifier for the IBM MQ libraries, the national language you want to use, the four-character IBM MQ queue manager name, and the destination for the job output.

3. Update RACF to allow CSQ4IVP1 to access its resources if IBM MQ security is active.

To run CSQ4IVP1 when IBM MQ security is enabled, you need a RACF user ID with authority to access the objects. For details of defining resources to RACF, see *Setting up security on z/OS*. The user ID that runs the IVP must have the following access authority:

Authority	Profile	Class
READ	ssid.DISPLAY.PROCESS	MQCMDS
UPDATE	ssid.SYSTEM.COMMAND.INPUT	MQQUEUE
UPDATE	ssid.SYSTEM.COMMAND.REPLY.MODEL	MQQUEUE
UPDATE	ssid.CSQ4IVP1.**	MQQUEUE
READ	ssid.BATCH	MQCONN

These requirements assume that all IBM MQ security is active. The RACF commands to activate IBM MQ security are shown in Figure 171 on page 1254. This example assumes that the queue manager name is CSQ1 and that the user ID of the person running sample CSQ4IVP1 is TS101.

```

RDEFINE MQCMDS CSQ1.DISPLAY.PROCESS
PERMIT CSQ1.DISPLAY.PROCESS CLASS(MQCMDS) ID(TS101) ACCESS(READ)

RDEFINE MQQUEUE CSQ1.SYSTEM.COMMAND.INPUT
PERMIT CSQ1.SYSTEM.COMMAND.INPUT CLASS(MQQUEUE) ID(TS101) ACCESS(UPDATE)

RDEFINE MQQUEUE CSQ1.SYSTEM.COMMAND.REPLY.MODEL
PERMIT CSQ1.SYSTEM.COMMAND.REPLY.MODEL CLASS(MQQUEUE) ID(TS101) ACCESS(UPDATE)

RDEFINE MQQUEUE CSQ1.CSQ4IVP1.**
PERMIT CSQ1.CSQ4IVP1.** CLASS(MQQUEUE) ID(TS101) ACCESS(UPDATE)

RDEFINE MQCONN CSQ1.BATCH
PERMIT CSQ1.BATCH CLASS(MQCONN) ID(TS101) ACCESS(READ)

```

Figure 171. RACF commands for CSQ4IVP1

## Running CSQ4IVP1

When you have completed these steps, start your queue manager. If the queue manager is already running and you have changed CSQINP2, you must stop the queue manager and restart it.

The IVP runs as a batch job. Customize the job card to meet the submission requirements of your installation.

## Checking the results of CSQ4IVP1

The IVP is split into 10 stages; each stage must complete with a zero completion code before the next stage is run. The IVP generates a report, listing:

- The name of queue manager that is being connected to.
- A one-line message showing the completion code and the reason code returned from each stage.
- A one-line informational message where appropriate.

A sample report is provided in Figure 172 on page 1256

For an explanation of the completion and reason codes, see the [../com.ibm.mq.ref.doc/q050270\\_.dita](#).

Some stages have more than one IBM MQ call and, in the event of failure, a message is issued indicating the specific IBM MQ call that returned the failure. Also, for some stages the IVP puts explanatory and diagnostic information into a comment field.

The IVP job requests exclusive control of certain queue manager objects and therefore should be single threaded through the system. However, there is no limit to the number of times the IVP can be run against your queue manager.

The functions performed by each stage are:

### Stage 1

Connect to the queue manager by issuing the **MQCONN** API call.

### Stage 2

Determine the name of the system-command input queue used by the command server to retrieve request messages. This queue receives display requests from Stage 5.

To do this, the sequence of calls is:

1. Issue an **MQOPEN** call, specifying the queue manager name, to open the queue manager object.
2. Issue an **MQINQ** call to find out the name of the system-command input queue.

3. Issue an **MQINQ** call to find out about various queue manager event switches.
4. Issue an **MQCLOSE** call to close the queue manager object.

On successful completion of this stage, the name of the system-command input queue is displayed in the comment field.

### Stage 3

Open an initiation queue using an **MQOPEN** call.

This queue is opened at this stage in anticipation of a trigger message, which arrives as a result of the command server replying to the request from Stage 5. The queue must be opened for input to meet the triggering criteria.

### Stage 4

Create a permanent dynamic queue using the CSQ4IVP1.MODEL queue as a model. The dynamic queue has the same attributes as the model from which it was created. This means that when the replies from the command server request in Stage 5 are written to this queue, a trigger message is written to the initiation queue opened in Stage 3.

Upon successful completion of this stage, the name of the permanent dynamic queue is indicated in the comment field.

### Stage 5

Issue an **MQPUT1** request to the command server command queue.

A message of type MQMT\_REQUEST is written to the system-command input queue requesting a display of process CSQ4IVP1. The message descriptor for the message specifies the permanent dynamic queue created in Stage 4 as the reply-to queue for the command server's response.

### Stage 6

Issue an **MQGET** request from the initiation queue. At this stage, a GET WAIT with an interval of 1 minute is issued against the initiation queue opened in Stage 3. The message returned is expected to be the trigger message generated by the command server's response messages being written to the reply-to queue.

### Stage 7

Delete the permanent dynamic queue created in Stage 4. As the queue still has messages on it, the MQCO\_PURGE\_DELETE option is used.

### Stage 8

1. Open a dynamic queue.
2. MQPUT a message with an expiry interval set.
3. Wait for the message to expire.
4. Attempt to MQGET the expired message.
5. MQCLOSE the queue.

### Stage 9

1. Open a dynamic queue.
2. MQPUT a message.
3. Issue MQCMIT to commit the current unit of work.
4. MQGET the message.
5. Issue MQBACK to backout the message.
6. MQGET the same message and ensure that the backout count is set to 1.
7. Issue MQCLOSE to close the queue.

### Stage 10

Disconnect from the queue manager using **MQDISC** .

After running the IVP, you can delete any objects that you no longer require.

If the IVP does not run successfully, try each step manually to find out which function is failing.

```
DATE : 2005.035 IBM MQ for z/OS - V6 PAGE : 0001
INSTALLATION VERIFICATION PROGRAM
PARAMETERS ACCEPTED. PROGRAM WILL CONNECT TO : CSQ1
,OBJECT QUALIFER : CSQ4IVP1
INSTALLATION VERIFICATION BEGINS :
STAGE 01 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR BRIDGE EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS EXCP FOR CHANNEL EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR SSL EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR INHIBITED EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR LOCAL EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR PERFORMANCE EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR REMOTE EVENTS
STAGE 02 INFO: QMGR EVENT SWITCH IS OFF FOR START/STOP EVENTS
STAGE 02 COMPLETE. COMPCODE : 0000 REASON CODE : 0000 SYSTEM.COMMAND.INPUT
STAGE 03 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
STAGE 04 COMPLETE. COMPCODE : 0000 REASON CODE : 0000 CSQ4IVP1.BAB9810EFEAC8980
STAGE 05 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
STAGE 06 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
STAGE 07 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
STAGE 08 COMPLETE. COMPCODE : 0000 REASON CODE : 0000 CSQ4IVP1.BAB9810F0070E645
STAGE 09 COMPLETE. COMPCODE : 0000 REASON CODE : 0000 CSQ4IVP1.BAB9812BA8706803
STAGE 10 COMPLETE. COMPCODE : 0000 REASON CODE : 0000
>>>>>>>> END OF REPORT <<<<<<<<<<<<
```

Figure 172. Sample report from CSQ4IVP1

## Testing for queue-sharing groups

The basic installation verification program CSQ4IVP1 tests non-shared queues.

CSQ4IVP1 can be used whether the queue manager is a member of a queue-sharing group or not. After running the basic IVP, you can test for shared queues by using the CSQ4IVP1 installation verification program with different queues. Also this tests that Db2 and the coupling facility are set up correctly.

## Preparing to run CSQ4IVP1 for a queue-sharing group

Before you run CSQ4IVP1:

1. Add the coupling facility structure that the IVP uses to your CFRM policy data set, as described in “Task 10: Set up the coupling facility” on page 1213. The supplied samples use a structure called APPLICATION1, but you can change this if you want.
2. Check that the IVP entries are in the CSQINP2 data set concatenation in the queue manager startup program. The IVP entries are supplied in member thlqual.SCSQPROC(CSQ4IVPG). If they are not, add the definitions supplied in thlqual.SCSQPROC(CSQ4IVPG) to your CSQINP2 concatenation. If the queue manager is currently running, you need to restart it so that these definitions can take effect.
3. Change the name of the coupling facility structure used in thlqual.SCSQPROC(CSQ4IVPG) if necessary.
4. The sample JCL, CSQ4IVPS, required to run the installation verification program for a queue-sharing group is in library thlqual.SCSQPROC.

Customize the CSQ4IVPS JCL with the high-level qualifier for the IBM MQ libraries, the national language you want to use, the four-character IBM MQ queue manager name, and the destination for the job output.

5. Update RACF to allow CSQ4IVP1 to access its resources if IBM MQ security is active.

To run CSQ4IVP1 when IBM MQ security is enabled, you need a RACF user ID with authority to access the objects. For details of defining resources to RACF, see Setting up security on z/OS . The user ID that runs the IVP must have the following access authority in addition to that required to run the basic IVP:



Authority	Profile	Class
UPDATE	ssid.CSQ4IVPG.**	MQQUEUE

These requirements assume that all IBM MQ security is active. The RACF commands to activate IBM MQ security are shown in Figure 173. This example assumes that the queue manager name is CSQ1 and that the user ID of the person running sample CSQ4IVP1 is TS101.

```
RDEFINE MQQUEUE CSQ1.CSQ4IVPG.**
PERMIT CSQ1.CSQ4IVPG.** CLASS(MQQUEUE) ID(TS101) ACCESS(UPDATE)
```

Figure 173. RACF commands for CSQ4IVP1 for a queue-sharing group

## Running CSQ4IVP1 for a queue-sharing group

When you have completed these steps, start your queue manager. If the queue manager is already running and you have changed CSQINP2, you must stop the queue manager and restart it.

The IVP runs as a batch job. Customize the job card to meet the submission requirements of your installation.

## Checking the results of CSQ4IVP1 for a queue-sharing group

The IVP for queue-sharing groups works in the same way as the basic IVP, except that the queues that are created are called CSQIVPG. xx. Follow the instructions given in “Checking the results of CSQ4IVP1” on page 1254 to check the results of the IVP for queue-sharing groups.

## Testing for distributed queuing

You can use the supplied installation verification program, CSQ4IVPX, to confirm that distributed queuing is operational.

## Overview of CSQ4IVPX job

CSQ4IVPX is a batch job that starts the channel initiator and issues the IBM MQ DISPLAY CHINIT command. This verifies that all major aspects of distributed queuing are operational, while avoiding the need to set up channel and network definitions.

## Preparing to run CSQ4IVPX

Before you run CSQ4IVPX:

1. The sample JCL, CSQ4IVPX, required to run the installation verification program is in library thlqual.SCSQPROC.

Customize the CSQ4IVPX JCL with the high-level qualifier for the IBM MQ libraries, the national language you want to use, the four-character queue manager name, and the destination for the job output.

2. Update RACF to allow CSQ4IVPX to access its resources if IBM MQ security is active. To run CSQ4IVPX when IBM MQ security is enabled, you need a RACF user ID with authority to access the objects. For details of defining resources to RACF, see *Setting up security on z/OS*. The user ID that runs the IVP must have the following access authority:

Authority	Profile	Class
CONTROL	ssid.START.CHINIT and ssid.STOP.CHINIT	MQCMDS
UPDATE	ssid.SYSTEM.COMMAND.INPUT	MQQUEUE
UPDATE	ssid.SYSTEM.CSQUTIL.*	MQQUEUE
READ	ssid.BATCH	MQCONN
READ	ssid.DISPLAY.CHINIT	MQCMDS

These requirements assume that the connection security profile ssid.CHIN has been defined (as shown in Connection security profiles for the channel initiator ), and that all IBM MQ security is active. The RACF commands to do this are shown in Figure 174 on page 1259. This example assumes that:

- The queue manager name is CSQ1
- The user ID of the person running sample CSQ4IVPX is TS101
- The channel initiator address space is running under the user ID CSQ1MSTR

3. Update RACF to allow the channel initiator address space the following access authority:

Authority	Profile	Class
READ	ssid.CHIN	MQCONN
UPDATE	ssid.SYSTEM.COMMAND.INPUT	MQQUEUE
UPDATE	ssid.SYSTEM.CHANNEL.INITQ	MQQUEUE
UPDATE	ssid.SYSTEM.CHANNEL.SYNCQ	MQQUEUE
ALTER	ssid.SYSTEM.CLUSTER.COMMAND.QUEUE	MQQUEUE
UPDATE	ssid.SYSTEM.CLUSTER.TRANSMIT.QUEUE	MQQUEUE
ALTER	ssid.SYSTEM.CLUSTER.REPOSITORY.QUEUE	MQQUEUE
CONTROL	ssid.CONTEXT.**	MQADMIN

The RACF commands to do this are also shown in Figure 174 on page 1259.

```

RDEFINE MQCMDS CSQ1.DISPLAY.DQM
PERMIT CSQ1.DISPLAY.DQM CLASS(MQCMDS) ID(TS101) ACCESS(READ)

RDEFINE MQCMDS CSQ1.START.CHINIT
PERMIT CSQ1.START.CHINIT CLASS(MQCMDS) ID(TS101) ACCESS(CONTROL)

RDEFINE MQCMDS CSQ1.STOP.CHINIT
PERMIT CSQ1.STOP.CHINIT CLASS(MQCMDS) ID(TS101) ACCESS(CONTROL)

RDEFINE MQQUEUE CSQ1.SYSTEM.COMMAND.INPUT
PERMIT CSQ1.SYSTEM.COMMAND.INPUT CLASS(MQQUEUE) ID(TS101,CSQ1MSTR) ACCESS(UPDATE)

RDEFINE MQQUEUE CSQ1.SYSTEM.CSQUTIL.*
PERMIT CSQ1.SYSTEM.CSQUTIL.* CLASS(MQQUEUE) ID(TS101) ACCESS(UPDATE)

RDEFINE MQCONN CSQ1.BATCH
PERMIT CSQ1.BATCH CLASS(MQCONN) ID(TS101) ACCESS(READ)

RDEFINE MQCONN CSQ1.CHIN
PERMIT CSQ1.CHIN CLASS(MQCONN) ID(CSQ1MSTR) ACCESS(READ)

RDEFINE MQQUEUE CSQ1.SYSTEM.CHANNEL.SYNCQ
PERMIT CSQ1.SYSTEM.CHANNEL.SYNCQ CLASS(MQQUEUE) ID(CSQ1MSTR) ACCESS(UPDATE)

RDEFINE MQQUEUE CSQ1.SYSTEM.CLUSTER.COMMAND.QUEUE
PERMIT CSQ1.SYSTEM.CLUSTER.COMMAND.QUEUE CLASS(MQQUEUE) ID(CSQ1MSTR) ACCESS(ALTER)

RDEFINE MQQUEUE CSQ1.SYSTEM.CLUSTER.TRANSMIT.QUEUE
PERMIT CSQ1.SYSTEM.CLUSTER.TRANSMIT.QUEUE CLASS(MQQUEUE) ID(CSQ1MSTR) ACCESS(UPDATE)

RDEFINE MQQUEUE CSQ1.SYSTEM.CLUSTER.REPOSITORY.QUEUE
PERMIT CSQ1.SYSTEM.CLUSTER.REPOSITORY.QUEUE CLASS(MQQUEUE) ID(CSQ1MSTR) ACCESS(ALTER)

RDEFINE MQQUEUE CSQ1.SYSTEM.CHANNEL.INITQ
PERMIT CSQ1.SYSTEM.CHANNEL.INITQ CLASS(MQQUEUE) ID(CSQ1MSTR) ACCESS(UPDATE)

RDEFINE MQADMIN CSQ1.CONTEXT.**
PERMIT CSQ1.CONTEXT.** CLASS(MQADMIN) ID(CSQ1MSTR) ACCESS(CONTROL)

```

Figure 174. RACF commands for CSQ4IVPX

## Running CSQ4IVPX

When you have completed these steps, start your queue manager.

The IVP runs as a batch job. Customize the job card to meet the submission requirements of your installation.

## Checking the results of CSQ4IVPX

CSQ4IVPX runs the CSQUTIL IBM MQ utility to issue three MQSC commands. The SYSPRINT output data set should look like Figure 175 on page 1260, although details might differ depending on your queue manager attributes.

- You should see the commands **(1)** each followed by several messages.
- The last message from each command should be “CSQ9022I ... NORMAL COMPLETION” **(2)**.
- The job as a whole should complete with return code zero **(3)**.

```

CSQU000I CSQUTIL IBM MQ for z/OS - V6
CSQU001I CSQUTIL Queue Manager Utility - 2005-05-09 09:06:48
COMMAND
CSQU127I CSQUTIL Executing COMMAND using input from CSQUCMD data set
CSQU120I CSQUTIL Connecting to queue manager CSQ1
CSQU121I CSQUTIL Connected to queue manager CSQ1
CSQU055I CSQUTIL Target queue manager is CSQ1
START CHINIT
(1)
CSQN205I COUNT= 2, RETURN=00000000, REASON=00000004
CSQM138I +CSQ1 CSQMSCHI CHANNEL INITIATOR STARTING
CSQN205I COUNT= 2, RETURN=00000000, REASON=00000000
CSQ9022I +CSQ1 CSQXCRPS ' START CHINIT' NORMAL COMPLETION
(2)
DISPLAY CHINIT
(1)
CSQN205I COUNT= 2, RETURN=00000000, REASON=00000004
CSQM137I +CSQ1 CSQMDDQM DISPLAY CHINIT COMMAND ACCEPTED
CSQN205I COUNT= 12, RETURN=00000000, REASON=00000000
CSQX830I +CSQ1 CSQXRDQM Channel initiator active
CSQX002I +CSQ1 CSQXRDQM Queue-sharing group is QSG1
CSQX831I +CSQ1 CSQXRDQM 8 adapter subtasks started, 8 requested
CSQX832I +CSQ1 CSQXRDQM 5 dispatchers started, 5 requested
CSQX833I +CSQ1 CSQXRDQM 0 SSL server subtasks started, 0 requested
CSQX840I +CSQ1 CSQXRDQM 0 channel connections current, maximum 200
CSQX841I +CSQ1 CSQXRDQM 0 channel connections active, maximum 200,
including 0 paused
CSQX842I +CSQ1 CSQXRDQM 0 channel connections starting,
0 stopped, 0 retrying
CSQX836I +CSQ1 Maximum channels - TCP/IP 200, LU 6.2 200
CSQX845I +CSQ1 CSQXRDQM TCP/IP system name is TCP/IP
CSQX848I +CSQ1 CSQXRDQM TCP/IP listener INDISP=QMGR not started
CSQX848I +CSQ1 CSQXRDQM TCP/IP listener INDISP=GROUP not started
CSQX849I +CSQ1 CSQXRDQM LU 6.2 listener INDISP=QMGR not started
CSQX849I +CSQ1 CSQXRDQM LU 6.2 listener INDISP=GROUP not started
CSQ9022I +CSQ1 CSQXCRPS ' DISPLAY CHINIT' NORMAL COMPLETION
(2)
STOP CHINIT
(1)
CSQN205I COUNT= 2, RETURN=00000000, REASON=00000004
CSQM137I +CSQ1 CSQMTCHI STOP CHINIT COMMAND ACCEPTED
CSQN205I COUNT= 2, RETURN=00000000, REASON=00000000
CSQ9022I +CSQ1 CSQXCRPS ' STOP CHINIT' NORMAL COMPLETION
(2)
CSQU057I CSQUCMDS 3 commands read
CSQU058I CSQUCMDS 3 commands issued and responses received, 0 failed
CSQU143I CSQUTIL 1 COMMAND statements attempted
CSQU144I CSQUTIL 1 COMMAND statements executed successfully
CSQU148I CSQUTIL Utility completed, return code=0
(3)

```

Figure 175. Example output from CSQ4IVPX

## Testing for C, C++, COBOL, PL/I, and CICS

You can test for C, C++, COBOL, PL/I, or CICS, using the sample applications supplied with IBM MQ.

The IVP (CSQ4IVP1) is supplied as a load module, and provides the samples as source modules. You can use these source modules to test different programming language environments.

For more information about sample applications, see *Sample programs for IBM MQ for z/OS* .

## Setting up communications with other queue managers

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

To define your distributed-queuing requirements, you need to define the following items:

- Define the channel initiator procedures and data sets
- Define the channel definitions
- Define the queues and other objects
- Define access security

To enable distributed queuing, you must perform the following three tasks:

- Customize the distributed queuing facility and define the IBM MQ objects required as described in *Defining system objects and "Preparing to customize your IBM MQ for z/OS queue managers"* on page 1194.
- Define access security as described in *Security considerations for the channel initiator on z/OS* .
- Set up your communications as described in *"Setting up communication for z/OS"* on page 1283.

If you are using queue-sharing groups, see *Distributed queuing and queue-sharing groups*.

See the following sections for additional considerations for using distributed queuing with IBM MQ for z/OS.

## Operator messages

Because the channel initiator uses a number of asynchronously operating dispatchers, operator messages might occur on the log out of chronological sequence.

## Channel operation commands

Channel operation commands generally involve two stages. When the command syntax has been checked and the existence of the channel verified, a request is sent to the channel initiator. Message CSQM134I or CSQM137I is sent to the command issuer to indicate the completion of the first stage. When the channel initiator has processed the command, further messages indicating its success or otherwise are sent to the command issuer along with message CSQ9022I or CSQ9023I. Any error messages generated could also be sent to the z/OS console.

All cluster commands except DISPLAY CLUSQMGR, however, work asynchronously. Commands that change object attributes update the object and send a request to the channel initiator. Commands for working with clusters are checked for syntax and a request is sent to the channel initiator. In both cases, message CSQM130I is sent to the command issuer indicating that a request has been sent. This message is followed by message CSQ9022I to indicate that the command has completed successfully, in that a request has been sent. It does not indicate that the cluster request has completed successfully. The requests sent to the channel initiator are processed asynchronously, along with cluster requests received from other members of the cluster. In some cases, these requests must be sent to the whole cluster to determine if they are successful or not. Any errors are reported to the z/OS on the system where the channel initiator is running. They are not sent to the command issuer.

## Undelivered-message queue

A Dead Letter handler is provided with IBM MQ for z/OS. See The dead-letter queue handler utility (CSQUDLQH) for more information.

## Queues in use

MCAs for receiver channels can keep the destination queues open even when messages are not being transmitted. This behavior results in the queues appearing to be 'in use'.

## Security changes

If you change security access for a user ID, the change might not take effect immediately. (See one of Security considerations for the channel initiator on z/OS , Profiles for queue security, and "Task 11: Implement your ESM security controls" on page 1214 for more information.)

## Communications stopped - TCP

If TCP is stopped for some reason and then restarted, the IBM MQ for z/OS TCP listener waiting on a TCP port is stopped.

Automatic channel-reconnect allows the channel initiator to detect that TCP/IP is unavailable and to automatically restart the TCP/IP listener when TCP/IP returns. This automatic restart alleviates the need for operations staff to notice the problem with TCP/IP and manually restart the listener. While the listener is out of action, the channel initiator can also be used to try the listener again at the interval specified by LSTRTMR in the channel initiator parameter module. These attempts can continue until TCP/IP returns and the listener successfully restarts automatically. For information about LSTRTMR, see Security concepts on z/OS and Distributed queuing messages (CSQX...).

## Communications stopped - LU6.2

If APPC is stopped, the listener is also stopped. Again, in this case, the listener automatically tries again at the LSTRTMR interval so that, if APPC restarts, the listener can restart too.

If the Db2 fails, shared channels that are already running continue to run, but any new channel start requests fail. When the Db2 is restored new requests are able to complete.

## z/OS Automatic Restart Management (ARM)

Automatic restart management (ARM) is a z/OS recovery function that can improve the availability of specific batch jobs or started tasks (for example, subsystems). It can therefore result in a faster resumption of productive work.

To use ARM, you must set up your queue managers and channel initiators in a particular way to make them restart automatically. For information, see Using the z/OS Automatic Restart Manager (ARM).

## Related concepts:

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Customizing IBM MQ for z/OS” on page 1198

Use this topic as a step by step guide for customizing your IBM MQ system.

“Monitoring and controlling channels on z/OS” on page 1264

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers.

“Setting up communication for z/OS” on page 1283

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. To succeed, it is necessary for the connection to be defined and available. This section explains how to define a connection.

“Preparing IBM MQ for z/OS for DQM with queue-sharing groups” on page 1288

Use the instructions in this section to configure distributed queuing with queue-sharing groups on IBM MQ for z/OS.

“Setting up communication for IBM MQ for z/OS using queue-sharing groups” on page 1293

When a distributed-queuing management channel is started, it attempts to use the connection specified in the channel definition. For this attempt to succeed, it is necessary for the connection to be defined and available.

## Defining IBM MQ objects

Use one of the IBM MQ command input methods to define IBM MQ objects. Refer to the information within this topic for further details about defining these objects.

Refer to “Monitoring and controlling channels on z/OS” on page 1264 for information about defining objects.

## Transmission queues and triggering channels

Define the following:

- A local queue with the usage of XMITQ for each sending message channel.
- Remote queue definitions.

A remote queue object has three distinct uses, depending upon the way the name and content are specified:


- Remote queue definition
- Queue manager alias definition
- Reply-to queue alias definition

These three ways are shown in Three ways of using the remote queue definition object.

Use the TRIGDATA field on the transmission queue to trigger the specified channel. For example:

```
DEFINE QLOCAL(MYXMITQ) USAGE(XMITQ) TRIGGER +
INITQ(SYSTEM.CHANNEL.INITQ) TRIGDATA(MYCHANNEL)
DEFINE CHL(MYCHANNEL) CHLTYPE(SDR) TRPTYPE(TCP) +
XMITQ(MYXMITQ) CONNAME('9.20.9.30(1555)')
```

The supplied sample CSQ4INXD gives additional examples of the necessary definitions.

 Loss of connectivity to the CF structure where the synchronization queue for shared channels is defined, or similar problems, might temporarily prevent a channel from starting. After problem resolution, if you are using a trigger type of FIRST and the channel fails to start when it is triggered, you must start the channel manually. If you want to automatically start triggered channels after problem resolution, consider setting the queue manager TRIGINT attribute to a value other than the default.

Setting the TRIGINT attribute to a value other than the default causes the channel initiator to retry starting the channel periodically while there are messages on the transmission queue.

## Synchronization queue

DQM requires a queue for use with sequence numbers and logical units of work identifiers (LUWID). You must ensure that a queue is available with the name SYSTEM.CHANNEL.SYNCQ (see Planning on z/OS ). This queue must be available otherwise the channel initiator cannot start.

Make sure that you define this queue using INDXTYPE(MSGID). This attribute improves the speed at which they can be accessed.

## Channel command queues

You need to ensure that a channel command queue exists for your system with the name SYSTEM.CHANNEL.INITQ.

If the channel initiator detects a problem with the SYSTEM.CHANNEL.INITQ, it is unable to continue normally until the problem is corrected. The problem could be one of the following:

- The queue is full
- The queue is not enabled for put
- The page set that the queue is on is full
- The channel initiator does not have the correct security authorization to the queue

If the definition of the queue is changed to GET(DISABLED) while the channel initiator is running, the initiator is unable to get messages from the queue, and terminates.

## Starting the channel initiator

Triggering is implemented using the channel initiator. On IBM MQ for z/OS, the initiator is started with the MQSC command START CHINIT.

## Stopping the channel initiator

The channel initiator is stopped automatically when you stop the queue manager. If you need to stop the channel initiator but not the queue manager, you can use the MQSC command STOP CHINIT.

## Monitoring and controlling channels on z/OS

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers.

Each z/OS queue manager has a DQM program (the *channel initiator* ) for controlling interconnections to remote queue managers using native z/OS facilities.

The implementation of these panels and commands on z/OS is integrated into the operations and control panels and the MQSC commands. No differentiation is made in the organization of these two sets of panels and commands.

You can also enter commands using Programmable Command Format (PCF) commands. See Automating administration tasks for information about using these commands.

The information in this section applies in all cases where the channel initiator is used for distributed queuing. It applies whether you are using queue-sharing groups, or intra-group queuing.



## The DQM channel control function

For an overview of the distributed queue management model, see “Message sending and receiving” on page 824.

The channel control function consists of panels, commands and programs, two synchronization queues, channel command queues, and the channel definitions. This topic is a brief description of the components of the channel control function.

- The channel definitions are held as objects in page set zero or in Db2, like other IBM MQ objects in z/OS.
- You use the operations and control panels, MQSC commands, or PCF commands to:
  - Create, copy, display, alter, and delete channel definitions
  - Start and stop channel initiators and listeners
  - Start, stop, and ping channels, reset channel sequence numbers, and resolve in-doubt messages when links cannot be re-established
  - Display status information about channels
  - Display information about DQM

In particular, you can use the CSQINPX initialization input data set to issue your MQSC commands. This set can be processed every time you start the channel initiator. For more information, see Initialization commands.

- There are two queues (SYSTEM.CHANNEL.SYNCQ and SYSTEM.QSG.CHANNEL.SYNCQ) used for channel re-synchronization purposes. Define these queues with INDXTYPE(MSGID) for performance reasons.
- The channel command queue (SYSTEM.CHANNEL.INITQ) is used to hold commands for channel initiators, channels, and listeners.
- The channel control function program runs in its own address space, separate from the queue manager, and comprises the channel initiator, listeners, MCAs, trigger monitor, and command handler.
- For queue-sharing groups and shared channels, see Shared queues and queue-sharing groups.
- For intra-group queuing, see Intra-group queuing

## Managing your channels on z/OS

Use the links in the following table for information about how to manage your channels, channel initiators, and listeners:

*Table 156. Channel tasks*

Task to be performed	MQSC command
Define a channel	DEFINE CHANNEL
Alter a channel definition	ALTER CHANNEL
Display a channel definition	DISPLAY CHANNEL
Delete a channel definition	DELETE CHANNEL
Start a channel initiator	START CHINIT
Stop a channel initiator	STOP CHINIT
Display channel initiator information	DISPLAY CHINIT
Start a channel listener	START LISTENER
Stop a channel listener	STOP LISTENER
Start a channel	START CHANNEL
Test a channel	PING CHANNEL

Table 156. Channel tasks (continued)

Task to be performed	MQSC command
Reset message sequence numbers for a channel	RESET CHANNEL
Resolve in-doubt messages on a channel	RESOLVE CHANNEL
Stop a channel	STOP CHANNEL
Display channel status	DISPLAY CHSTATUS
Display cluster channels	DISPLAY CLUSQMGR

**Related concepts:**

“Using the panels and the commands”

You can use the MQSC commands, the PCF commands, or the operations and control panels to manage DQM.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

“Customizing IBM MQ for z/OS” on page 1198

Use this topic as a step by step guide for customizing your IBM MQ system.

“Setting up communication for z/OS” on page 1283

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. To succeed, it is necessary for the connection to be defined and available. This section explains how to define a connection.

“Preparing IBM MQ for z/OS for DQM with queue-sharing groups” on page 1288

Use the instructions in this section to configure distributed queuing with queue-sharing groups on IBM MQ for z/OS.

“Setting up communication for IBM MQ for z/OS using queue-sharing groups” on page 1293

When a distributed-queuing management channel is started, it attempts to use the connection specified in the channel definition. For this attempt to succeed, it is necessary for the connection to be defined and available.

**Using the panels and the commands:**

You can use the MQSC commands, the PCF commands, or the operations and control panels to manage DQM.

For information about the syntax of the MQSC commands, see Script (MQSC) Commands. For information about PCF commands, see Introduction to Programmable Command Formats.

**Using the initial panel**

For an introduction to invoking the operations and control panels, using the function keys, and getting help, see Administering IBM MQ for z/OS.

**Note:** To use the operations and control panels, you must have the correct security authorization; see Administering IBM MQ for z/OS and sub topics for more information. Figure 176 on page 1267 shows the panel that is displayed when you start a panel session. The text after the panel explains the actions you perform in this panel.

```

IBM MQ for z/OS - Main Menu

Complete fields. Then press Enter.

Action 1 0. List with filter 4. Manage
1. List or Display 5. Perform
2. Define like 6. Start
3. Alter 7. Stop
8. Command
Object type CHANNEL +
Name *
Disposition A Q=Qmgr, C=Copy, P=Private, G=Group,
S=Shared, A=All

Connect name MQ25 - local queue manager or group
Target queue manager . . . MQ25
- connected or remote queue manager for command input
Action queue manager . . . MQ25 - command scope in group
Response wait time 10 5 - 999 seconds

(C) Copyright IBM Corporation 1993,2005. All rights reserved.

Command ==> _____
F1=Help F2=Split F3=Exit F4=Prompt F9=SwapNext F10=Messages
F12=Cancel

```

Figure 176. The operations and controls initial panel

From this panel, you can:

- Select the action you want to perform by typing in the appropriate number in the **Action** field.
- Specify the object type that you want to work with. Press F4 for a list of object types if you are not sure what they are.
- Display a list of objects of the type specified. Type in an asterisk (\*) in the **Name** field and press enter to display a list of objects (of the type specified) that have already been defined on this subsystem. You can then select one or more objects to work with in sequence. Figure 177 on page 1268 shows a list of channels produced in this way.
- Specify the disposition in the queue-sharing group of the objects you want to work with in the **Disposition** field. The disposition determines where the object is kept and how the object behaves.
- Choose the local queue manager, or queue-sharing group to which you want to connect in the **Connect name** field. If you want the commands to be issued on a remote queue manager, choose either the **Target queue manager** field or the **Action queue manager** field, depending upon whether the remote queue manager is not or is a member of a queue-sharing group. If the remote queue manager is not a member of a queue-sharing group, choose the **Target queue manager** field. If the remote queue manager is a member of a queue-sharing group, choose the **Action queue manager** field.
- Choose the wait time for responses to be received in the **Response wait time** field.

```

List Channels - MQ25 Row 1 of 8

Type action codes, then press Enter. Press F11 to display connection status.
1=Display 2=Define like 3=Alter 4=Manage 5=Perform
6=Start 7=Stop

Name Type Disposition Status
<> * CHANNEL ALL MQ25
- SYSTEM.DEF.CLNTCONN CLNTCONN QMGR MQ25
- SYSTEM.DEF.CLUSRCVR CLUSRCVR QMGR MQ25 INACTIVE
- SYSTEM.DEF.CLUSSDR CLUSSDR QMGR MQ25 INACTIVE
- SYSTEM.DEF.RECEIVER RECEIVER QMGR MQ25 INACTIVE
- SYSTEM.DEF.REQUESTER REQUESTER QMGR MQ25 INACTIVE
- SYSTEM.DEF.SENDER SENDER QMGR MQ25 INACTIVE
- SYSTEM.DEF.SERVER SERVER QMGR MQ25 INACTIVE
- SYSTEM.DEF.SVRCONN SVRCONN QMGR MQ25 INACTIVE
***** End of list *****

Command ==> _____
F1=Help F2=Split F3=Exit F4=Filter F5=Refresh F7=Bkwd
F8=Fwd F9=SwapNext F10=Messages F11=Status F12=Cancel

```

Figure 177. Listing channels

### Defining a channel:

You can define a channel using MQSC commands or using the operations and control panels.

To define a channel using the MQSC commands, use DEFINE CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	2 (Define like)
Object type	channel type (for example SENDER) or CHANNEL
Name	
Disposition	The location of the new object.

You are presented with some panels to complete with information about the name and attributes you want for the channel you are defining. They are initialized with the default attribute values. Change any you want before pressing enter.

**Note:** If you entered CHANNEL in the **object type** field, you are presented with the Select a Valid Channel Type panel first.

If you want to define a channel with the same attributes as an existing channel, put the name of the channel you want to copy in the **Name** field on the initial panel. The panels are initialized with the attributes of the existing object.

For information about the channel attributes, see Channel attributes

### Note:

1. Name all the channels in your network uniquely. As shown in Network diagram showing all channels, including the source and target queue manager names in the channel name is a good way to do this naming.

After you have defined your channel you must secure your channel, see “Securing a channel” on page 1271

### **Altering a channel definition:**

You can alter a channel definition using MQSC commands or using the operations and control panels.

To alter a channel definition using the MQSC commands, use ALTER CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

<b>Field</b>	<b>Value</b>
Action	3 (Alter)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.ALTER
Disposition	The location of the stored object.

You are presented with some panels containing information about the current attributes of the channel. Change any of the unprotected fields that you want by over typing the new value, and then press enter to change the channel definition.

For information about the channel attributes, see Channel attributes.

### **Displaying a channel definition:**

You can display a channel definition using MQSC commands or using the operations and control panels.

To display a channel definition using the MQSC commands, use DISPLAY CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

<b>Field</b>	<b>Value</b>
Action	1 (List or Display)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.DISPLAY
Disposition	The location of the object.

You are presented with some panels displaying information about the current attributes of the channel.

For information about the channel attributes, see Channel attributes.

### Deleting a channel definition:

You can delete a channel definition using MQSC commands or using the operations and control panels.

To delete a channel definition using the MQSC commands, use DELETE CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	4 (Manage)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.DELETE
Disposition	The location of the object.

You are presented with another panel. Select function type 1 on this panel.

Press enter to delete the channel definition; you are asked to confirm that you want to delete the channel definition by pressing enter again.

**Note:** The channel initiator has to be running before a channel definition can be deleted (except for client-connection channels).

### Displaying information about the channel initiator:

You can display information about the channel initiator using MQSC commands or using the operations and control panels.

To display information about the channel initiator using the MQSC commands, use DISPLAY CHINIT.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	1 (Display)
Object type	SYSTEM
Name	Blank

You are presented with another panel. Select function type 1 on this panel.

#### **Note:**

1. Displaying distributed queuing information might take some time if you have lots of channels.
2. The channel initiator has to be running before you can display information about distributed queuing.

## Securing a channel:

You can secure a channel using MQSC commands or using the operations and control panels.

To secure a channel using the MQSC commands, use SET CHLAUTH.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	8

You are presented with an editor within which you can provide an MQSC command, in this case a CHLAUTH command, see Figure 178. When you have finished typing in the command, the plus signs (+) are needed. Type PF3 to exit from the editor and submit the command to the command server.

```
***** Top of Data *****
000001 SET CHLAUTH(SYSTEM.DEF.SVRCONN) +
000002 TYPE(SSLPEERMAP) +
000003 SSLPEER('CN="John Smith"') +
000004 MCAUSER('PUBLIC')
***** Bottom of Data *****

Command ==> Scroll ==> PAGE
F1=Help F3=Exit F4=LineEdit F12=Cancel
```

Figure 178. Command Entry

The output of the command is then presented to you, see Figure 179

```
***** Top of Data *****
000001 CSQU000I CSQUTIL IBM MQ for z/OS V7.1.0
000002 CSQU001I CSQUTIL Queue Manager Utility - 2011-04-20 14:42:58
000003 COMMAND TGTQMGR(MQ23) RESPTIME(30)
000004 CSQU127I Executing COMMAND using input from CSQUCMD data set
000005 CSQU120I Connecting to MQ23
000006 CSQU121I Connected to queue manager MQ23
000007 CSQU055I Target queue manager is MQ23
000008 SET CHLAUTH(SYSTEM.DEF.SVRCONN) +
000009 TYPE(SSLPEERMAP) +
000010 SSLPEER('CN="John Smith"') +
000011 MCAUSER('PUBLIC')
000012 CSQN205I COUNT= 2, RETURN=00000000, REASON=00000000
000013 CSQ9022I !MQ23 CSQMCA ' SET CHLAUTH' NORMAL COMPLETION
000014 CSQU057I 1 commands read
000015 CSQU058I 1 commands issued and responses received, 0 failed
000016 CSQU143I 1 COMMAND statements attempted
000017 CSQU144I 1 COMMAND statements executed successfully
000018 CSQU148I CSQUTIL Utility completed, return code=0
Command ==> Scroll ==> PAGE
F1=Help F3=Exit F5=Rfind F6=Rchange F9=SwapNext F12=Cancel
```

Figure 179. Command Output

## Starting a channel initiator:

You can start a channel initiator using MQSC commands or using the operations and control panels.

To start a channel initiator using the MQSC commands, use START CHINIT.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	6 (Start)
Object type	SYSTEM
Name	Blank

The Start a System Function panel is displayed. The text following the following panel explains what action to take:

```
Start a System Function

Select function type, complete fields, then press Enter to start system
function.

Function type _ 1. Channel initiator
2. Channel listener
Action queue manager . . . : MQ25

Channel initiator
JCL substitution _____

Channel listener
Inbound disposition . . . Q G=Group, Q=Qmgr
Transport type _ L=LU6.2, T=TCP/IP
LU name (LU6.2) _____
Port number (TCP/IP) . . . 1414
IP address (TCP/IP) _____

Command ==> _____
F1=Help F2=Split F3=Exit F9=SwapNext F10=Messages F12=Cancel
```

Figure 180. Starting a system function

Select function type 1 (channel initiator), and press enter.



### Stopping a channel initiator:

You can stop a channel initiator using MQSC commands or using the operations and control panels.

To stop a channel initiator using the MQSC commands, use STOP CHINIT.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	7 (Stop)
Object type	SYSTEM
Name	Blank

The Stop a System Function panel is displayed. The text following the panel explains how you to use this panel:

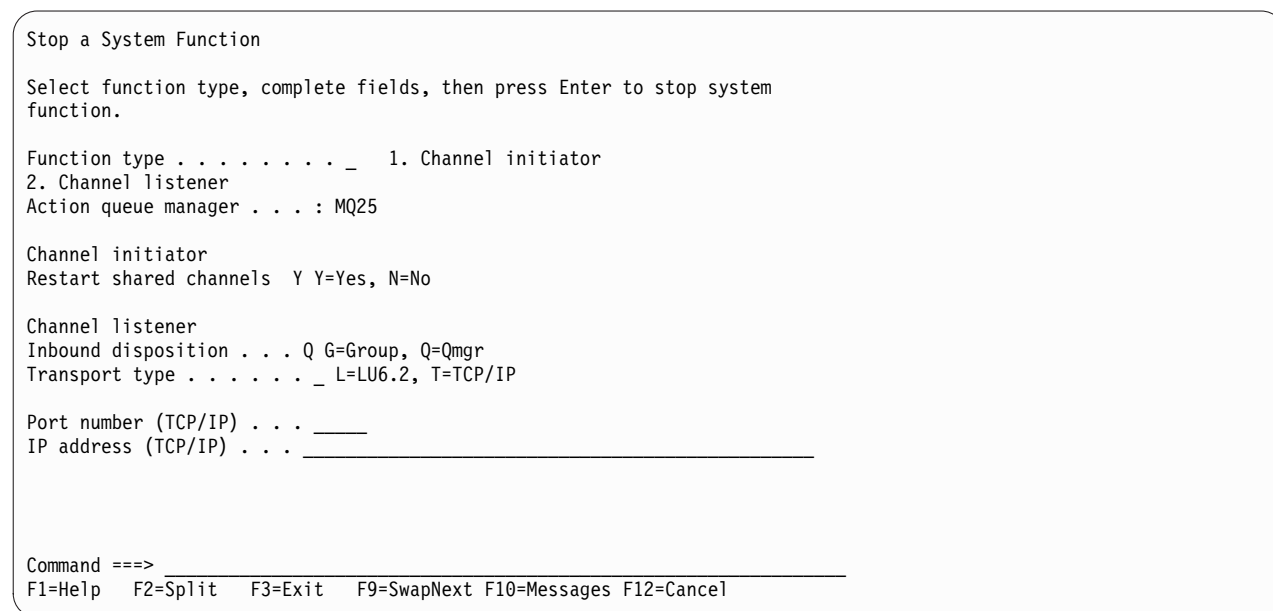


Figure 181. Stopping a function control

Select function type 1 (channel initiator) and press enter.

The channel initiator waits for all running channels to stop in quiesce mode before it stops.

**Note:** If some of the channels are receiver or requester channels that are running but not active, a stop request issued to either the receiver or sender channel initiator causes it to stop immediately.

However, if messages are flowing, the channel initiator waits for the current batch of messages to complete before it stops.

### Starting a channel listener:

You can start a channel listener using MQSC commands or using the operations and control panels.

To start a channel listener using the MQSC commands, use START LISTENER.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	6 (Start)
Object type	SYSTEM
Name	Blank

The Start a System Function panel is displayed (see Figure 180 on page 1272 ).

Select function type 2 (channel listener). Select Inbound disposition. Select Transport type. If the Transport type is L, select LU name. If the Transport type is T, select Port number and (optionally) IP address. Press enter.

**Note:** For the TCP/IP listener, you can start multiple combinations of Port and IP address.

### Stopping a channel listener:

You can stop a channel listener using MQSC commands or using the operations and control panels.

To stop a channel listener using the MQSC commands, use STOP LISTENER.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	7 (Stop)
Object type	SYSTEM
Name	Blank

The Stop a System Function panel is displayed (see Figure 181 on page 1273 ).

Select function type 2 (channel listener). Select Inbound disposition. Select Transport type. If the transport type is "T", select Port number and (optionally) IP address. Press enter.

**Note:** For a TCP/IP listener, you can stop specific combinations of Port and IP address, or you can stop all combinations.

### Starting a channel:

You can start a channel using MQSC commands or using the operations and control panels.

To start a channel using the MQSC commands, use START CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	6 (Start)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.USE
Disposition	The disposition of the object.

The Start a Channel panel is displayed. The text following the panel explains how to use the panel:




Figure 182. Starting a channel

Select the disposition of the channel instance and on which queue manager it is to be started.

Press enter to start the channel.

## Starting a shared channel:

To start a shared channel, and keep it on a nominated channel initiator, use disposition = S (on the START CHANNEL command, specify CHLDISP(FIXSHARED)).

There can be only one instance of the shared channel running at a time. Attempts to start a second instance of the channel fail.

When you start a channel in this way, the following rules apply to that channel:

- You can stop the channel from any queue manager in the queue-sharing group. You can stop it even if the channel initiator on which it was started is not running at the time you issue the stop-channel request. When the channel has stopped, you can restart it by specifying disposition = S (CHLDISP(FIXSHARED)) on the same, or another, channel initiator. You can also start it by specifying disposition = A (CHLDISP(SHARED)).
- If the channel is in the starting or retry state, you can restart it by specifying disposition = S (CHLDISP(FIXSHARED)) on the same or a different channel initiator. You can also start it by specifying disposition = A (CHLDISP(SHARED)).
- The channel is eligible to be trigger started when it goes into the inactive state. Shared channels that are trigger started always have a shared disposition (CHLDISP(SHARED)).
- The channel is eligible to be started with CHLDISP(FIXSHARED), on any channel initiator, when it goes into the inactive state. You can also start it by specifying disposition = A (CHLDISP(SHARED)).
- The channel is not recovered by any other active channel initiator in the queue-sharing group when the channel initiator on which it was started is stopped with SHARED(RESTART), or when the channel initiator terminates abnormally. The channel is recovered only when the channel initiator on which it was started is next restarted. This stops failed channel-recovery attempts being passed to other channel initiators in the queue-sharing group, which would add to their workload.

## Testing a channel:

You can test a channel using MQSC commands or using the operations and control panels.

To test a channel using the MQSC commands, use PING CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	5 (Perform)
Object type	SENDER, SERVER, or CHANNEL
Name	CHANNEL.TO.USE
Disposition	The disposition of the channel object.

The Perform a Channel Function panel is displayed. The text following the panel explains how to use the panel:

Perform a Channel Function

Select function type, complete fields, then press Enter.

Function type . . . . . \_ 1. Reset 3. Resolve with commit  
2. Ping 4. Resolve with backout

Channel name . . . . . : CHANNEL.TO.USE  
Channel type . . . . . : SENDER  
Description . . . . . : Description of CHANNEL.TO.USE

Disposition . . . . . P P=Private on MQ25  
S=Shared on MQ25  
A=Shared on any queue manager

Sequence number for reset . . 1 1 - 999999999  
Data length for ping . . . . 16 16 - 32768

Command ==> \_\_\_\_\_  
F1=Help F2=Split F3=Exit F9=SwapNext F10=Messages F12=Cancel

Figure 183. Testing a channel

Select function type 2 (ping).

Select the disposition of the channel for which the test is to be done and on which queue manager it is to be tested.

The data length is initially set to 16. Change it if you want and press enter.

**Resetting message sequence numbers for a channel:**

You can reset message sequence numbers for a channel using MQSC commands or using the operations and control panels.

To reset channel sequence numbers using the MQSC commands, use RESET CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	5 (Perform)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.USE
Disposition	The disposition of the channel object.

The Perform a Channel Function panel is displayed (see Figure 183 ).

Select Function type 1 (reset).

Select the disposition of the channel for which the reset is to be done and on which queue manager it is to be done.

The **sequence number** field is initially set to one. Change this value if you want, and press enter.

### Resolving in-doubt messages on a channel:

You can resolve in-doubt messages on a channel using MQSC commands or using the operations and control panels.

To resolve in-doubt messages on a channel using the MQSC commands, use RESOLVE CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	5 (Perform)
Object type	SENDER, SERVER, or CHANNEL
Name	CHANNEL.TO.USE
Disposition	The disposition of the object.

The Perform a Channel Function panel is displayed (see Figure 183 on page 1277 ).

Select Function type 3 or 4 (resolve with commit or backout). (See "In-doubt channels" on page 844 for more information.)

Select the disposition of the channel for which resolution is to be done and which queue manager it is to be done on. Press enter.

### Stopping a channel:

You can stop a channel using MQSC commands or using the operations and control panels.

To stop a channel using the MQSC commands, use STOP CHANNEL.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	7 (Stop)
Object type	channel type (for example SENDER) or CHANNEL
Name	CHANNEL.TO.USE
Disposition	The disposition of the object.

The Stop a Channel panel is displayed. The text following the panel explains how to use the panel:

### Stop a Channel

Complete fields, then press Enter to stop channel.

Channel name . . . . . : CHANNEL.TO.USE  
Channel type . . . . . : SENDER  
Description . . . . . : Description of CHANNEL.TO.USE

Disposition . . . . . P P=Private on MQ25  
A=Shared on any queue manager

Stop mode . . . . . 1 1. Quiesce 2. Force  
Stop status . . . . . 1 1. Stopped 2. Inactive

Queue manager . . . . . \_\_\_\_\_  
Connection name . . . . . \_\_\_\_\_

Command ==> \_\_\_\_\_  
F1=Help F2=Split F3=Exit F9=SwapNext F10=Messages F12=Cancel

Figure 184. Stopping a channel

Select the disposition of the channel for which the stop is to be done and on which queue manager it is to be stopped.

Choose the stop mode that you require:

#### Quiesce

The channel stops when the current message is completed and the batch is then ended, even if the batch size value has not been reached and there are messages already waiting on the transmission queue. No new batches are started. This mode is the default.

**Force** The channel stops immediately. If a batch of messages is in progress, an 'in-doubt' situation can result.

Choose the queue manager and connection name for the channel you want to stop.

Choose the status that you require:

#### Stopped

The channel is not restarted automatically, and must be restarted manually. This mode is the default if no queue manager or connection name is specified. If a name is specified, it is not allowed.

#### Inactive

The channel is restarted automatically when required. This mode is the default if a queue manager or connection name is specified.

Press enter to stop the channel.

See “Stopping and quiescing channels” on page 842 for more information. For information about restarting stopped channels, see “Restarting stopped channels” on page 844.

**Note:** If a shared channel is in a retry state and the channel initiator on which it was started is not running, a STOP request for the channel is issued on the queue manager where the command was entered.

## Displaying channel status:

You can display channel status by using MQSC commands, or by using the operations and control panels.

To display the status of a channel or a set of channels using the MQSC commands, use DISPLAY CHSTATUS.

**Note:** Displaying channel status information can take some time if you have lots of channels.

Using the operations and control panels on the List Channel panel (see Figure 177 on page 1268 ), a summary of the channel status is shown for each channel as follows:

INACTIVE	No connections are active
<i>status</i>	One connection is active
<i>nmn status</i>	More than one connection is current and all current connections have the same status
<i>nmn CURRENT</i>	More than one connection is current and the current connections do not all have the same status
Blank	IBM MQ is unable to determine how many connections are active (for example, because the channel initiator is not running)
	<b>Note:</b> For channel objects with the disposition GROUP, no status is displayed.

where *nmn* is the number of active connections, and *status* is one of the following:

INIT	INITIALIZING
BIND	BINDING
START	STARTING
RUN	RUNNING
STOP	STOPPING or STOPPED
RETRY	RETRYING
REQST	REQUESTING

To display more information about the channel status, press the Status key (F11) on the List Channel or the Display, or Alter channel panels to display the List Channels - Current Status panel (see Figure 185 on page 1281 ).



List Channels - Current Status - MQ25 Row 1 of 16

Type action codes, then press Enter. Press F11 to display saved status.  
1=Display current status

Channel name	Connection name	State
Start time	Messages Last message time	Type Disposition
<> *	CHANNEL ALL	MQ25
_	RMA0.CIRCUIT.ACL.F RMA1	STOP
_	2005-03-21 10.22.36 557735	2005-03-24 09.51.11 SENDER PRIVATE MQ25
_	RMA0.CIRCUIT.ACL.N RMA1	
_	2005-03-21 10.23.09 378675	2005-03-24 09.51.10 SENDER PRIVATE MQ25
_	RMA0.CIRCUIT.CL.F RMA2	
_	2005-03-24 01.12.51 45544	2005-03-24 09.51.08 SENDER PRIVATE MQ25
_	RMA0.CIRCUIT.CL.N RMA2	
_	2005-03-24 01.13.55 45560	2005-03-24 09.51.11 SENDER PRIVATE MQ25
_	RMA1.CIRCUIT.CL.F RMA1	
_	2005-03-21 10.24.12 360757	2005-03-24 09.51.11 RECEIVER PRIVATE MQ25
_	RMA1.CIRCUIT.CL.N RMA1	
_	2005-03-21 10.23.40 302870	2005-03-24 09.51.09 RECEIVER PRIVATE MQ25
***** End of list *****		
Command ==>		
F1=Help F2=Split F3=Exit F4=Filter F5=Refresh F7=Bkwd		
F8=Fwd F9=SwapNext F10=Messages F11=Saved F12=Cancel		

Figure 185. Listing channel connections

The values for status are as follows:

INIT	INITIALIZING
BIND	BINDING
START	STARTING
RUN	RUNNING
STOP	STOPPING or STOPPED
RETRY	RETRYING
REQST	REQUESTING
DOUBT	STOPPED and INDOUBT(YES)

See "Channel states" on page 834 for more information.

You can press F11 to see a similar list of channel connections with saved status; press F11 to get back to the current list. The saved status does not apply until at least one batch of messages has been transmitted on the channel.

Use action code 1 or a slash (/) to select a connection and press enter. The Display Channel Connection Current Status panels are displayed.

## Displaying cluster channels:

You can display cluster channels using MQSC commands or using the operations and control panels.

To display all the cluster channels that have been defined (explicitly or using auto-definition), use the MQSC command, DISPLAY CLUSQMGR.

Using the operations and control panels, starting from the initial panel, complete these fields and press enter:

Field	Value
Action	1 (List or Display)
Object type	CLUSCHL
Name	*

You are presented with a panel like figure Figure 186, in which the information for each cluster channel occupies three lines, and includes its channel, cluster, and queue manager names. For cluster-sender channels, the overall state is shown.

```
List Cluster-queue-manager Channels - MQ25 Row 1 of 9

Type action codes, then press Enter. Press F11 to display connection status.
1=Display 5=Perform 6=Start 7=Stop

Channel name Connection name State
Type Cluster name Suspended
Cluster queue manager name Disposition
<> * - MQ25
- TO.MQ90.T HURSLEY.MACH90.COM(1590)
- CLUSRCVR VJH01T - MQ25 N
- MQ90 - MQ25
- TO.MQ95.T HURSLEY.MACH95.COM(1595) RUN
- CLUSSDRA VJH01T - MQ25 N
- MQ95 - MQ25
- TO.MQ96.T HURSLEY.MACH96.COM(1596) RUN
- CLUSSDRB VJH01T - MQ25 N
- MQ96 - MQ25
***** End of list *****

Command ==> _____
F1=Help F2=Split F3=Exit F4=Filter F5=Refresh F7=Bkwd
F8=Fwd F9=SwapNext F10=Messages F11=Status F12=Cancel
```

Figure 186. Listing cluster channels

To display full information about one or more channels, type Action code 1 against their names and press enter. Use Action codes 5, 6, or 7 to perform functions (such as ping, resolve, and reset), and start or stop a cluster channel.

To display more information about the channel status, press the Status key (F11).

## Preparing IBM MQ for z/OS to use the zEnterprise Data Compression Express facility:

The zEnterprise® Data Compression (zEDC) Express facility is available for certain models of IBM Z machines, using a minimum z/OS level of z/OS Version 2.1.

See IBM zEnterprise Data Compression (zEDC) for z/OS for further information.

To start using zEDC for channel compression, you must:

- Configure the zEDC Software License in the IFAPRDxx parmlib member.
- Ensure that the channel initiator user ID has READ authority to the FPZ.ACCELERATOR.COMPRESSION profile in the RACF FACILITY CLASS, or the equivalent in the external security manager (ESM) that your enterprise uses.
- Configure the channel with COMPMSG(ZLIBFAST) at both the sending and receiving ends.

Once configured, zlib compression is used to compress and uncompress messages flowing across the channel.

Compression is performed in the zEDC when the size of data to be compressed is above a size defined by the zEDC Express facility, currently 4 KB. For messages below the threshold size, compression is performed in the software.

## Setting up communication for z/OS

When a distributed-queuing management channel is started, it tries to use the connection specified in the channel definition. To succeed, it is necessary for the connection to be defined and available. This section explains how to define a connection.

DQM is a remote queuing facility for IBM MQ. It provides channel control programs for the queue manager that form the interface to communication links. These links are controllable by the system operator. The channel definitions held by distributed queuing management use these connections.

Choose from one of the two forms of communication protocol that can be used for z/OS:

- “Defining a TCP connection on z/OS” on page 1284
- “Defining an LU6.2 connection for z/OS using APPC/MVS” on page 1287

Each channel definition must specify only one protocol as the transmission protocol (Transport Type) attribute. A queue manager can use more than one protocol to communicate.

You might also find it helpful to refer to Example configuration - IBM MQ for z/OS . If you are using queue sharing groups, see “Setting up communication for IBM MQ for z/OS using queue-sharing groups” on page 1293.

### Related concepts:

“Using the panels and the commands” on page 1266

You can use the MQSC commands, the PCF commands, or the operations and control panels to manage DQM.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

“Customizing IBM MQ for z/OS” on page 1198

Use this topic as a step by step guide for customizing your IBM MQ system.

“Monitoring and controlling channels on z/OS” on page 1264

Use the DQM commands and panels to create, monitor, and control the channels to remote queue managers.

“Preparing IBM MQ for z/OS for DQM with queue-sharing groups” on page 1288

Use the instructions in this section to configure distributed queuing with queue-sharing groups on IBM MQ for z/OS.

“Setting up communication for IBM MQ for z/OS using queue-sharing groups” on page 1293

When a distributed-queuing management channel is started, it attempts to use the connection specified in the channel definition. For this attempt to succeed, it is necessary for the connection to be defined and available.

### Defining a TCP connection on z/OS:

To define a TCP connection, there are a number of settings to configure.

The TCP address space name must be specified in the TCP system parameters data set, *tcPIP.TCPIP.DATA*. In the data set, a “*TCPIPJOBNAME TCPIP\_proc*” statement must be included.

If you are using a firewall, you need to configure allow connections from the channel initiator to the addresses in the channels, and from the remote connections into the queue manager.

Typically the definition for a firewall configures the sending IP address and port to the destination IP address and port:

- A z/OS image can have more than one host name, and you might need to configure the firewall with multiple host addresses as the source address.  
You can use the NETSTAT HOME command to display these names and addresses.
- A channel initiator can have multiple listeners on different ports, so you need to configure these ports.
- If you are using a shared port for a queue sharing group you must configure the shared port as well.

The channel initiator address space must have authority to read the data set. The following techniques can be used to access your *TCPIP.DATA* data set, depending on which TCP/IP product and interface you are using:

- Environment variable, *RESOLVER\_CONFIG*
- HFS file, */etc/resolv.conf*
- *//SYSTCPD DD* statement
- *//SYSTCPDD DD* statement
- *jobname/userid.TCPIP.DATA*
- *SYS1.TCPPARMS(TCPDATA)*
- *zapname.TCPIP.DATA*

You must also be careful to specify the high-level qualifier for TCP/IP correctly.

You need a suitably configured Domain Name System (DNS) server, capable of both Name to IP address translation and IP address to Name translation.

**Note:** Some changes to the resolver configuration require a recycle of applications using it, for example, IBM MQ.

For more information, see the following:

- Base TCP/IP system
- z/OS UNIX System Services.

Each TCP channel when started uses TCP resources; you might need to adjust the following parameters in your PROFILE.TCPIP configuration data set:

**ACBPOOLSIZE**

Add one per started TCP channel, plus one

**CCBPOOLSIZE**

Add one per started TCP channel, plus one per DQM dispatcher, plus one

**DATABUFFERPOOLSIZE**

Add two per started TCP channel, plus one

**MAXFILEPROC**

Controls how many channels each dispatcher in the channel initiator can handle.

This parameter is specified in the BPXPRMxx member of SYSI.PARMLIB. Ensure that you specify a value large enough for your needs.

By default, the channel initiator is only capable of binding to IP addresses associated with the stack named in the TCPNAME queue manager attribute. To allow the channel initiator to communicate using additional TCP/IP stacks on the system, change the TCPSTACK queue manager attribute to MULTIPLE.

**Related concepts:**

“Sending end”

At the sending end of the TCP/IP connection, there are a number of settings to configure.

“Receiving on TCP” on page 1286

At the receiving end of the TCP/IP connection, there are a number of settings to configure.

“Using the TCP listener backlog option” on page 1286

When receiving on TCP/IP, a maximum number of outstanding connection requests is set. These outstanding requests can be considered a *backlog* of requests waiting on the TCP/IP port for the listener to accept the request.

*Sending end:*

At the sending end of the TCP/IP connection, there are a number of settings to configure.

The connection name (CONNNAME) field in the channel definition must be set to either the host name (for example MVSHUR1) or the TCP network address of the target. The TCP network address can be in IPv4 dotted decimal form (for example 127.0.0.1) or IPv6 hexadecimal form (for example 2001:DB8:0:0:0:0:0:0). If the connection name is a host name, a TCP name server is required to convert the host name into a TCP host address. (This requirement is a function of TCP, not IBM MQ.)

On the initiating end of a connection (sender, requester, and server channel types) it is possible to provide an optional port number for the connection, for example:

**Connection name**

192.0.2.0(1555)

In this case the initiating end attempts to connect to a receiving program listening on port 1555.

**Note:** The default port number of 1414 is used if an optional port number is not specified.

The channel initiator can use any TCP/IP stack which is active and available. By default, the channel initiator binds its outbound channels to the default IP address for the TCP/IP stack named in the TCPNAME queue manager attribute. To connect through a different stack, you need to specify either the host name or IP address of the stack in the LOCLADDR attribute of the channel.

*Receiving on TCP:*

At the receiving end of the TCP/IP connection, there are a number of settings to configure.

Receiving channel programs are started in response to a startup request from the sending channel. To do so, a listener program has to be started to detect incoming network requests and start the associated channel. You start this listener program with the START LISTENER command, or using the operations and control panels.

By default:

- The TCP Listener program uses port 1414 and listens on all addresses available to your TCP stack.
- TCP/IP listeners can bind only to addresses associated with the TCP/IP stack named in the TCPNAME queue manager attribute.

To start listeners for other addresses, or all available TCP stacks, set your TCPSTACK queue manager attribute to 'MULTIPLE'.

You can start your TCP listener program to listen only on a specific address or host name by specifying IPADDR in the START LISTENER command. For more information, see Listeners.

*Using the TCP listener backlog option:*

When receiving on TCP/IP, a maximum number of outstanding connection requests is set. These outstanding requests can be considered a *backlog* of requests waiting on the TCP/IP port for the listener to accept the request.

The default listener backlog value on z/OS is 255. If the backlog reaches this values, the TCP/IP connection is rejected and the channel is not able to start.

For MCA channels, this results in the channel going into a RETRY state and retrying the connection at a later time.

For client connections, the client receives an MQRC\_Q\_MGR\_NOT\_AVAILABLE reason code from MQCONN and can retry the connection at a later time.

## Defining an LU6.2 connection for z/OS using APPC/MVS:

To define an LU6.2 connection there are a number of settings to configure.

### APPC/MVS setup

Each instance of the channel initiator must have the name of the LU that it is to use defined to APPC/MVS, in the APPCPMxx member of SYS1.PARMLIB, as in the following example:

```
LUADD ACBNAME(luname) NOSCHED TPDATA(CSQ.APPCTP)
```

*luname* is the name of the logical unit to be used. NOSCHED is required; TPDATA is not used. No additions are necessary to the ASCHPMxx member, or to the APPC/MVS TP profile data set.

The side information data set must be extended to define the connections used by DQM. See the supplied sample CSQ4SIDE for details of how to do this using the APPC utility program ATBSDFMU. For details of the TPNAME values to use, see the *Multiplatform APPC Configuration Guide* ( "Redbook" ) and the following table for information:

Table 157. Settings on the local z/OS system for a remote queue manager platform

Remote platform	TPNAME
z/OS or MVS	The same as TPNAME in the corresponding side information about the remote queue manager.
IBM i	The same as the compare value in the routing entry on the IBM i system.
HP Integrity NonStop Server	The same as the TPNAME specified in the receiver-channel definition.
UNIX and Linux systems	The same as TPNAME in the corresponding side information about the remote queue manager.
Windows	As specified in the Windows Run Listener command, or the invocable Transaction Program that was defined using TpSetup on Windows.

If you have more than one queue manager on the same machine, ensure that the TPnames in the channel definitions are unique.

See the *Multiplatform APPC Configuration Guide* also for information about the VTAM definitions that might be required.

In an environment where the queue manager is communicating using APPC with a queue manager on the same or another z/OS system, ensure that either the VTAM definition for the communicating LU specifies SECACPT(ALREADYV), or that there is a RACF APPCLU profile for the connection between LUs, which specifies CONVSEC(ALREADYV).

The z/OS command VARY ACTIVE must be issued against both base and listener LUs before attempting to start either inbound or outbound communications.

**Related concepts:**

“Connecting to LU 6.2”

To connect to LU 6.2, there are a number of settings to configure.

“Receiving on LU 6.2”

To receive on LU 6.2, there are a number of settings to configure.

*Connecting to LU 6.2:*

To connect to LU 6.2, there are a number of settings to configure.

The connection name (CONNNAME) field in the channel definition must be set to the symbolic destination name, as specified in the side information data set for APPC/MVS.

The LU name to use (defined to APPC/MVS as described previously) must also be specified in the channel initiator parameters. It must be set to the same LU that is used for receiving by the listener.

The channel initiator uses the “SECURITY(SAME)” APPC/MVS option, so it is the user ID of the channel initiator address space that is used for outbound transmissions, and is presented to the receiver.

*Receiving on LU 6.2:*

To receive on LU 6.2, there are a number of settings to configure.

Receiving MCAs are started in response to a startup request from the sending channel. To do so, a listener program has to be started to detect incoming network requests and start the associated channel. The listener program is an APPC/MVS server. You start it with the START LISTENER command, or using the operations and control panels. You must specify the LU name to use with a symbolic destination name defined in the side information data set. The local LU so identified must be the same as the one used for outbound transmissions, as set in the channel initiator parameters.

**Preparing IBM MQ for z/OS for DQM with queue-sharing groups**

Use the instructions in this section to configure distributed queuing with queue-sharing groups on IBM MQ for z/OS.

For an example configuration using queue-sharing groups, see Example configuration - IBM MQ for z/OS using queue-sharing groups. For a message channel planning example using queue-sharing groups, see Message channel planning example for z/OS using queue-sharing groups.

You need to create and configure the following components to enable distributed queuing with queue-sharing groups:

- LU 6.2 and TCP/IP listeners
- Transmission queues and triggering
- Message channel agents
- Synchronization queue

After you have created the components you need to set up the communication, see “Setting up communication for IBM MQ for z/OS using queue-sharing groups” on page 1293.

For information about how to monitor and control channels when using queue-sharing groups, see “Monitoring and controlling channels on z/OS” on page 1264.

See the following sections for queue-sharing group concepts and benefits.



## Class of service

A shared queue is a type of local queue that offers a different class of service. Messages on a shared queue are stored in a coupling facility (CF), which allows them to be accessed by all queue managers in the queue-sharing group. A message on a shared queue must be a message of length no more than 100 MB.

## Generic interface

A queue-sharing group has a generic interface that allows the network to view the group as a single entity. This view is achieved by having a single generic address that can be used to connect to any queue manager within the group.

Each queue manager in the queue-sharing group listens for inbound session requests on an address that is logically related to the generic address. For more information see “LU 6.2 and TCP/IP listeners for queue-sharing groups” on page 1290.

## Load-balanced channel start

A shared transmission queue can be serviced by an outbound channel running on any channel initiator in the queue-sharing group. Load-balanced channel start determines where a start channel command is targeted. An appropriate channel initiator is chosen that has access to the necessary communications subsystem. For example, a channel defined with TRPTYPE(LU6.2) cannot be started on a channel initiator that only has access to a TCP/IP subsystem.

The choice of channel initiator is dependent on the channel load and the headroom of the channel initiator. The channel load is the number of active channels as a percentage of the maximum number of active channels allowed as defined in the channel initiator parameters. The headroom is the difference between the number of active channels and the maximum number allowed.

Inbound shared channels can be load-balanced across the queue-sharing group by use of a generic address, as described in “LU 6.2 and TCP/IP listeners for queue-sharing groups” on page 1290.

## Shared channel recovery

The following table shows the types of shared-channel failure and how each type is handled.

Type of failure:	What happens:
Channel initiator communications subsystem failure	The channels dependent on the communications subsystem enter channel retry, and are restarted on an appropriate queue-sharing group channel initiator by a load-balanced start command.
Channel initiator failure	The channel initiator fails, but the associated queue manager remains active. The queue manager monitors the failure and initiates recovery processing.
Queue manager failure	The queue manager fails (failing the associated channel initiator). Other queue managers in the queue-sharing group monitor the event and initiate peer recovery.
Shared status failure	Channel state information is stored in Db2, so a loss of connectivity to Db2 becomes a failure when a channel state change occurs. Running channels can carry on running without access to these resources. On a failed access to Db2, the channel enters retry.

Shared channel recovery processing on behalf of a failed system requires connectivity to Db2 to be available on the system managing the recovery to retrieve the shared channel status.

## Client channels

Client connection channels can benefit from the high availability of messages in queue-sharing groups that are connected to the generic interface instead of being connected to a specific queue manager. For

more information, see Client connection channels.

**Related concepts:**

“Customizing IBM MQ for z/OS” on page 1198

Use this topic as a step by step guide for customizing your IBM MQ system.

“Configuring distributed queuing” on page 802

This section provides more detailed information about intercommunication between IBM MQ installations, including queue definition, channel definition, triggering, and sync point procedures

“Setting up communications with other queue managers” on page 1261

This section describes the IBM MQ for z/OS preparations you need to make before you can start to use distributed queuing.

“Clusters and queue-sharing groups” on page 1292

You can make your shared queue available to a cluster in a single definition. To do so you specify the name of the cluster when you define the shared queue.

“Channels and serialization” on page 1292

During shared queue peer recovery, message channel agents that process messages on shared queues serialize their access to the queues.

**Related information:**

Shared queues and queue-sharing groups

Intra-group queuing

**LU 6.2 and TCP/IP listeners for queue-sharing groups:**

The group LU 6.2 and TCP/IP listeners listen on an address that is logically connected to the generic address.

For the LU 6.2 listener, the specified LUGROUP is mapped to the VTAM generic resource associated with the queue-sharing group. For an example of setting up this technology, see “Defining an LU6.2 connection for z/OS using APPC/MVS” on page 1287.

For the TCP/IP listener, the specified port can be connected to the generic address in one of the following ways:

- For a front-end router such as the IBM Network Dispatcher, inbound connect requests are forwarded from the router to the members of the queue-sharing group.
- For TCP/IP Sysplex Distributor, each listener that is running and is listening on a particular address that is set up as a Distributed DVIPA is allocated a proportion of the incoming requests. For an example of setting up this technology, see Using Sysplex Distributor

## Transmission queues and triggering for queue-sharing groups: z/OS

A shared transmission queue is used to store messages before they are moved from the queue-sharing group to the destination.

It is a shared queue and it is accessible to all queue managers in the queue-sharing group.

### Triggering

A triggered shared queue can generate more than one trigger message for a satisfied trigger condition. There is one trigger message generated for each local initiation queue defined on a queue manager in the queue-sharing group associated with the triggered shared queue.

For distributed queuing, each channel initiator receives a trigger message for a satisfied shared transmission queue trigger condition. However, only one channel initiator actually processes the triggered start, and the others fail safely. The triggered channel is then started with a load balanced start (see “Preparing IBM MQ for z/OS for DQM with queue-sharing groups” on page 1288 ) that is triggered to start channel QSG.TO.QM2. To create a shared transmission queue, use the IBM MQ commands (MQSC) as shown in the following example:

```
DEFINE QLOCAL(QM2) DESCR('Transmission queue to QM2') +
USAGE(XMITQ) QSGDISP(SHARED) +
CFSTRUCT(APPLICATION1) INITQ(SYSTEM.CHANNEL.INITQ) +
TRIGGER TRIGDATA(QSG.TO.QM2)
```

### Message channel agents for queue-sharing groups:

A channel can only be started on a channel initiator if it has access to a channel definition for a channel with that name.

A message channel agent is an IBM MQ program that controls the sending and receiving of messages. Message channel agents move messages from one queue manager to another; there is one message channel agent at each end of a channel.

A channel definition can be defined to be private to a queue manager or stored on the shared repository and available anywhere (a group definition). This means that a group defined channel is available on any channel initiator in the queue-sharing group.

**Note:** The private copy of the group definition can be changed or deleted.

To create group channel definitions, use the IBM MQ commands (MQSC) as shown in the following examples:

```
DEFINE CHL(QSG.TO.QM2) CHLTYPE(SDR) +
TRPTYPE(TCP) CONNAME(QM2.MACH.IBM.COM) +
XMITQ(QM2) QSGDISP(GROUP)
DEFINE CHL(QM2.TO.QSG) CHLTYPE(RCVR) TRPTYPE(TCP) +
QSGDISP(GROUP)
```

There are two perspectives from which to look at the message channel agents used for distributed queuing with queue-sharing groups:

### Inbound

An inbound channel is a shared channel if it is connected to the queue manager through the group listener. It is connected either through the generic interface to the queue-sharing group, then directed to a queue manager within the group, or targeted at the group port of a specific queue manager or the luname used by the group listener.

## Outbound

An outbound channel is a shared channel if it moves messages from a shared transmission queue. In the example commands, sender channel QSG.T0.QM2 is a shared channel because its transmission queue, QM2 is defined with QSGDISP(SHARED).

### Synchronization queue for queue-sharing groups:

Shared channels have their own shared synchronization queue called SYSTEM.QSG.CHANNEL.SYNCQ.

This synchronization queue is accessible to any member of the queue-sharing group. (Private channels continue to use the private synchronization queue. See “Defining IBM MQ objects” on page 1263 ). This means that the channel can be restarted on a different queue manager and channel initiator instance within the queue-sharing group in the event of failure of the communications subsystem, channel initiator, or queue manager. For further information, see “Preparing IBM MQ for z/OS for DQM with queue-sharing groups” on page 1288.

DQM with queue-sharing groups requires that a shared queue is available with the name SYSTEM.QSG.CHANNEL.SYNCQ. This queue must be available so that a group listener can successfully start.

If a group listener fails because the queue was not available, the queue can be defined and the listener can be restarted without recycling the channel initiator. The non-shared channels are not affected.

Make sure that you define this queue using INDXTYPE(MSGID). This definition improves the speed at which messages on the queue can be accessed.

### Clusters and queue-sharing groups:

You can make your shared queue available to a cluster in a single definition. To do so you specify the name of the cluster when you define the shared queue.

Users in the network see the shared queue as being hosted by each queue manager within the queue-sharing group. (The shared queue is not advertised as being hosted by the queue-sharing group). Clients can start sessions with all members of the queue-sharing group to put messages to the same shared queue.

For more information, see “Configuring a queue manager cluster” on page 902.

### Channels and serialization:

During shared queue peer recovery, message channel agents that process messages on shared queues serialize their access to the queues.

If a queue manager in a queue-sharing group fails while a message channel agent is dealing with uncommitted messages on one or more shared queues, the channel and the associated channel initiator will end, and shared queue peer recovery will take place for the queue manager.

Because shared queue peer recovery is an asynchronous activity, peer channel recovery might try to simultaneously restart the channel in another part of the queue sharing group before shared queue peer recovery is complete. If this event happens, committed messages might be processed ahead of the messages still being recovered. To ensure that messages are not processed out of sequence in this way, message channel agents that process messages on shared queues serialize their access to these queues.

An attempt to start a channel for which shared queue peer recovery is still in progress might result in a failure. An error message indicating that recovery is in progress is issued, and the channel is put into retry state. Once queue manager peer recovery is complete, the channel can restart at the time of the next retry.

An attempt to RESOLVE, PING, or DELETE a channel can fail for the same reason.

### **Setting up communication for IBM MQ for z/OS using queue-sharing groups:**

When a distributed-queuing management channel is started, it attempts to use the connection specified in the channel definition. For this attempt to succeed, it is necessary for the connection to be defined and available.

Choose from one of the two forms of communication protocol that can be used:

- TCP
- LU 6.2 through APPC/MVS

You might find it useful to refer to Example configuration - IBM MQ for z/OS using queue-sharing groups.

*Defining a TCP connection for queue-sharing groups:*

To define a TCP connection for a queue-sharing group, certain attributes on the sending and receiving end must be configured.

For information about setting up your TCP, see “Defining a TCP connection on z/OS” on page 1284.

### **Sending end**

The connection name (CONNNAME) field in the channel definition to connect to your queue-sharing group must be set to the generic interface of your queue-sharing group (see Queue-sharing groups ). For more details, refer to Using Sysplex Distributor.

### **Receiving on TCP using a queue-sharing group**

Receiving shared channel programs are started in response to a startup request from the sending channel. To do so, a listener must be started to detect incoming network requests and start the associated channel. You start this listener program with the START LISTENER command, using the inbound disposition of the group, or using the operations and control panels.

All group listeners in the queue-sharing group must be listening on the same port. If you have more than one channel initiator running on a single MVS image you can define virtual IP addresses and start your TCP listener program to only listen on a specific address or host name by specifying IPADDR in the START LISTENER command. (For more information, see START LISTENER.)

*Defining an LU 6.2 connection on z/OS:*

To define an LU 6.2 connection for a queue-sharing group, certain attributes on the sending and receiving end must be configured.

For information about setting up APPC/MVS, see *Setting up communication for z/OS* .

### **Connecting to APPC/MVS (LU 6.2)**

The connection name (CONNNAME) field in the channel definition to connect to your queue-sharing group must be set to the symbolic destination name, as specified in the side information data set for APPC/MVS. The partner LU specified in this symbolic destination must be the generic resource name. For more details, see *Defining yourself to the network using generic resources*.

### **Receiving on LU 6.2 using a generic interface**

Receiving shared MCAs are started in response to a startup request from the sending channel. To do so, a group listener program must be started to detect incoming network requests and start the associated channel. The listener program is an APPC/MVS server. You start it with the START LISTENER command, using an inbound disposition group, or using the operations and control panels. You must specify the LU name to use a symbolic destination name defined in the side information data set. For more details, see *Defining yourself to the network using generic resources*.

## **Using IBM MQ with IMS**

The IBM MQ -IMS adapter, and the IBM MQ - IMS bridge are the two components which allow IBM MQ to interact with IMS.

To configure IBM MQ and IMS to work together, you must complete the following tasks:

- “Setting up the IMS adapter”
- “Setting up the IMS bridge” on page 1302

### **Related concepts:**

“Configuring queue managers on z/OS” on page 1194

Use these instructions to configure queue managers on IBM MQ for z/OS.

“Using IBM MQ with CICS” on page 1303

To use IBM MQ with CICS, you must configure the IBM MQ CICS adapter and, optionally, the IBM MQ CICS bridge components.

“Using OTMA exits in IMS” on page 1305

Use this topic if you want to use IMS Open Transaction Manager Access exits with IBM MQ for z/OS.

### **Related reference:**

“Upgrading and applying service to Language Environment or z/OS Callable Services” on page 1303

The actions you must take vary according to whether you use CALLLIBS or LINK, and your version of SMP/E.

### **Related information:**

IBM MQ and IMS

IMS and IMS bridge applications on IBM MQ for z/OS

## **Setting up the IMS adapter**

To use IBM MQ within IMS requires the IBM MQ - IMS adapter (generally referred to as the IMS adapter).

This topic tells you how to make the IMS adapter available to your IMS subsystem. If you are not familiar with tailoring an IMS subsystem, see the *IMS Knowledge Center*.

To make the IMS adapter available to IMS applications, follow these steps:

1. Define IBM MQ to IMS as an external subsystem using the IMS external subsystem attach facility (ESAF).

See “Defining IBM MQ to IMS” on page 1296.

2. Include the IBM MQ load library thlqual.SCSQAUTH in the JOBLIB or STEPLIB concatenation in the JCL for your IMS control region and for any dependent region that connects to IBM MQ (if it is not in the LPA or link list). If your JOBLIB or STEPLIB is not authorized, also include it in the DFSESL concatenation after the library containing the IMS modules (usually IMS RESLIB).

Also include thlqual.SCSQANLx (where x is the language letter).

If DFSESL is present, then SCSQAUTH and SCSQANLx need to be included in the concatenation or added to LNKLIST. Adding to the STEPLIB or JOBLIB concatenation in the JCL is not sufficient.

3. Copy the IBM MQ assembler program CSQQDEFV from thlqual.SCSQASMS to a user library.
4. The supplied program, CSQQDEFV, contains one subsystem name CSQ1 identified as default with an IMS language interface token (LIT) of MQM1. You can retain this name for testing and installation verification.

For production subsystems, you change the NAME=CSQ1 to your own subsystem name, or use CSQ1. You can add further subsystem definitions as required. See “Defining IBM MQ queue managers to the IMS adapter” on page 1300 for further information on LITs.

5. Assemble and link-edit the program to produce the CSQQDEFV load module. For the assembly, include the library thlqual.SCSQMACS in your SYSLIB concatenation; use the link-edit parameter RENT. This is shown in the sample JCL in thlqual.SCSQPROC(CSQ4DEFV).
6. Include the user library containing the module CSQQDEFV that you created in the JOBLIB or STEPLIB concatenation in the JCL for any dependent region that connects to IBM MQ. Put this library before the SCSQAUTH because SCSQAUTH has a default load module. If you do not do this, you will receive a user 3041 abend from IMS.
7. If the IMS adapter detects an unexpected IBM MQ error, it issues a z/OS SNAP dump to DD name CSQSNAP and issues reason code MQRC\_UNEXPECTED\_ERROR to the application. If the CSQSNAP DD statement was not in the IMS dependent region JCL, no dump is taken. If this happens, you could include the CSQSNAP DD statement in the JCL and rerun the application. However, because some problems might be intermittent, it is recommended that you include the CSQSNAP DD statement to capture the reason for failure at the time it occurs.
8. If you want to use dynamic IBM MQ calls (described in Dynamically calling the IBM MQ stub ), build the dynamic stub, as shown in Figure 187 on page 1296.
9. If you want to use the IMS trigger monitor, define the IMS trigger monitor application CSQQTRMN, and perform PSBGEN and ACBGEN. See “Setting up the IMS trigger monitor” on page 1301.
10. If you are using RACF to protect resources in the OPERCMDS class, ensure that the userid associated with your IBM MQ queue manager address space has authority to issue the MODIFY command to any IMS system to which it might connect.

```

//DYNSTUB EXEC PGM=IEWL,PARM='RENT,REUS,MAP,XREF'
//SYSPRINT DD SYSOUT=*
//ACSQMOD DD DISP=SHR,DSN=thlqua1.SCSQLOAD
//IMSLIB DD DISP=SHR,DSN=ims.reslib
//SYSLMOD DD DISP=SHR,DSN=private.load1
//SYSUT1 DD UNIT=SYSDA,SPACE=(CYL,1)
//SYSLIN DD *
 INCLUDE ACSQMOD(CSQQSTUB)
 INCLUDE IMSLIB(DFSLI000)
 ALIAS MQCONN,MQCONN,MQDISC MQI entry points
 ALIAS MQGET,MQPUT,MQPUT1 MQI entry points
 ALIAS MQOPEN,MQCLOSE MQI entry points
 ALIAS MQBACK,MQCMIT MQI entry points
 ALIAS CSQBBAK,CSQBCMT MQI entry points
 ALIAS MQINQ,MQSET MQI entry points
 ALIAS DFSPLI,PLITDLI IMS entry points
 ALIAS DFSCOBOL,CBLTDLI IMS entry points
 ALIAS DFSFOR,FORTDLI IMS entry points
 ALIAS DFSASM,ASMTDLI IMS entry points
 ALIAS DFSPASCL,PASTDLI IMS entry points
 ALIAS DFHEI01,DFHEI1 IMS entry points
 ALIAS DFSAIBLI,AIBTDLI IMS entry points
 ALIAS DFSESS,DSNCLI,DSNHCLI IMS entry points
 ALIAS MQCRTMH,MQDLTMH,MQDLTMP IMS entry points
 ALIAS MQINQMP,MQSETMP,MQMHBUFF,MQBUFMH IMS entry points
 MODE AMODE(31),RMODE(24) Note RMODE setting
 NAME CSQQDYS(R)
/*
1Specify the name of a library accessible to IMS applications that
want to make dynamic calls to WebSphere MQ.

```

Figure 187. Sample JCL to link-edit the dynamic call stub. This includes the IMS language interface module and the IBM MQ IMS stub CSQQSTUB.

#### Related concepts:

“Setting up the IMS bridge” on page 1302

The IBM MQ - IMS bridge is an optional component that enables IBM MQ to input and output to and from existing programs and transactions that are not IBM MQ-enabled.

#### Related information:

IBM MQ and IMS

IMS and IMS bridge applications on IBM MQ for z/OS

#### Defining IBM MQ to IMS:

IBM MQ must be defined to the IMS control region, and to each dependent region accessing that IBM MQ queue manager. To do this, you must create a subsystem member (SSM) in the IMS.PROCLIB library, and identify the SSM to the applicable IMS regions.

#### Placing the subsystem member entry in IMS.PROCLIB

Each SSM entry in IMS.PROCLIB defines a connection from an IMS region to a different queue manager.

To name an SSM member, concatenate the value (one to four alphanumeric characters) of the IMSID field of the IMS IMSCTRL macro with any name (one to four alphanumeric characters) defined by your site.

One SSM member can be shared by all the IMS regions, or a specific member can be defined for each region. This member contains as many entries as there are connections to external subsystems. Each entry is an 80-character record.



## Positional parameters

The fields in this entry are:

SSN,LIT,ESMT,RTT,REO,CRC
--------------------------

where:

### SSN

Specifies the IBM MQ queue manager name. It is required, and must contain one through four characters.

### LIT

Specifies the language interface token (LIT) supplied to IMS. This field is required, its value must match one in the CSQQDEFV module.

### ESMT

Specifies the external subsystem module table (ESMT). This table specifies which attachment modules must be loaded by IMS. CSQQESMT is the required value for this field.

### RTT

This option is not supported by IBM MQ.

### REO

Specifies the region error option (REO) to be used if an IMS application references a non-operational external subsystem or if resources are unavailable at create thread time. This field is optional and contains a single character, which can be:

**R** Passes a return code to the application, indicating that the request for IBM MQ services failed.

**Q** Ends the application with an abend code U3051, backs out activity to the last commit point, does a PSTOP of the transaction, and requeues the input message. This option only applies when an IMS application tries to reference a non-operational external subsystem or if the resources are unavailable at create thread time.

IBM MQ completion and reason codes are returned to the application if the IBM MQ problem occurs while IBM MQ is processing the request; that is, after the adapter has passed the request on to IBM MQ.

**A** Ends the application with an abend code of U3047 and discards the input message. This option only applies when an IMS application references a non-operational external subsystem or if the resources are unavailable at create thread time.

IBM MQ completion and reason codes are returned to the application if the IBM MQ problem occurs while IBM MQ is processing the request; that is, after the adapter has passed the request on to IBM MQ.

### CRC

This option can be specified but is not used by IBM MQ.

An example SSM entry is:

```
CSQ1, MQM1, CSQQESMT, , R,
```

where:

<b>CSQ1</b>	The default subsystem name as supplied with IBM MQ. You can change this to suit your installation.
<b>MQM1</b>	The default LIT as supplied in CSQQDEFV.
<b>CSQQESMT</b>	The external subsystem module name. You must use this value.
<b>R</b>	REO option.

### Keyword parameters

IBM MQ parameters can be specified in keyword format; to do this you must specify SST=DB2. Other parameters are as described in Positional parameters, and shown in the following example:

```
SST=DB2, SSN=SYS3, LIT=MQM3, ESMT=CSQQESMT
```

where:

<b>SYS3</b>	The subsystem name
<b>MQM3</b>	The LIT as supplied in CSQQDEFV
<b>CSQQESMT</b>	The external subsystem module name

### Specifying the SSM EXEC parameter

Specify the SSM EXEC parameter in the startup procedure of the IMS control region. This parameter specifies the one-character to four-character subsystem member name (SSM).

If you specify the SSM for the IMS control region, any dependent region running under the control region can attach to the IBM MQ queue manager named in the IMS.PROCLIB member specified by the SSM parameter. The IMS.PROCLIB member name is the IMS ID (IMSID= *xxxx*) concatenated with the one to four characters specified in the SSM EXEC parameter. The IMS ID is the IMSID parameter of the IMSCTRL generation macro.

IMS lets you define as many external subsystem connections as are required. More than one connection can be defined for different IBM MQ queue managers. All IBM MQ connections must be within the same z/OS system. For a dependent region, you can specify a dependent region SSM or use the one specified for the control region. You can specify different region error options (REOs) in the dependent region SSM member and the control region SSM member. Table 158 on page 1299 shows the different possibilities of SSM specifications.

Table 158. SSM specifications options

SSM for control region	SSM for dependent region	Action	Comments
No	No	None	No external subsystem can be connected.
No	Yes	None	No external subsystem can be connected.
Yes	No	Use the control region SSM	Applications scheduled in the region can access external subsystems identified in the control region SSM. Exits and control blocks for each attachment are loaded into the control region and the dependent region address spaces.
Yes	Yes (empty)	No SSM is used for the dependent region	Applications scheduled in this region can access DL/I databases only. Exits and control blocks for each attachment are loaded into the control region address space.
Yes	Yes (not empty)	Check the dependent region SSM with the control region SSM	Applications scheduled in this region can access only external subsystems identified in both SSMs. Exits and control blocks for each attachment are loaded into the control region and the dependent region address spaces.

There is no specific parameter to control the maximum number of SSM specification possibilities.

### Preloading the IMS adapter

The performance of the IMS adapter can be improved if it is preloaded by IMS. Preloading is controlled by the DFSMPLxx member of IMS.PROCLIB: see "IMS Administration Guide: System" for more information. The IBM MQ module names to specify are:

CSQACLST	CSQAMLST	CSQAPRH	CSQAVICM	CSQFSALM	CSQQDEFV
CSQQCONN	CSQQDISC	CSQQTERM	CSQQINIT	CSQQBACK	CSQQCMMT
CSQQESMT	CSQQPREP	CSQQTTHD	CSQQWAIT	CSQQNORM	CSQQSSOF
CSQQSSON	CSQFSTAB	CSQQRESV	CSQQSNOP	CSQQCMND	CSQQCVER
CSQQTMID	CSQQTRGI	CSQQCON2	CSQBAPPL		

**V8.0.0.4** For more information on the use of IBM MQ classes for JMS, see Using IBM MQ classes for JMS in IMS.

Current releases of IMS support preloading IBM MQ modules from PDS-E format libraries in MPP and BMP regions only. Any other type of IMS region does not support preloading from PDS-E libraries. If preloading is required for any other type of region, then the IBM MQ modules that are provided must be copied to a PDS format library.

## Defining IBM MQ queue managers to the IMS adapter:

The names of the IBM MQ queue managers and their corresponding language interface tokens (LITs) must be defined in the queue manager definition table.

Use the supplied CSQQDEFX macro to create the CSQQDEFV load module. Figure 188 shows the syntax of this assembler macro.

```
CSQQDEFX TYPE=ENTRY|DEFAULT,NAME=qmgr-name,LIT=token
or
CSQQDEFX TYPE=END
```

Figure 188. CSQQDEFX macro syntax

### Parameters

#### **TYPE=ENTRY|DEFAULT**

Specify either TYPE=ENTRY or TYPE=DEFAULT as follows:

##### **TYPE=ENTRY**

Specifies that a table entry describing an IBM MQ queue manager available to an IMS application is to be generated. If this is the first entry, the table header is also generated, including a CSQQDEFV CSECT statement.

##### **TYPE=DEFAULT**

As for TYPE=ENTRY. The queue manager specified is the default queue manager to be used when **MQCONN** or **MQCONNX** specifies a name that is all blanks. There must be only one such entry in the table.

##### **NAME= *qmgr-name***

Specifies the name of the queue manager, as specified with **MQCONN** or **MQCONNX**.

##### **LIT= *token***

Specifies the name of the language interface token (LIT) that IMS uses to identify the queue manager.

An **MQCONN** or **MQCONNX** call associates the *name* input parameter and the *hconn* output parameter with the name label and, therefore, the LIT in the CSQQDEFV entry. Further IBM MQ calls passing the *hconn* parameter use the LIT from the CSQQDEFV entry identified in the **MQCONN** or **MQCONNX** call to direct calls to the IBM MQ queue manager defined in the IMS SSM PROCLIB member with that same LIT.

In summary, the *name* parameter on the **MQCONN** or **MQCONNX** call identifies a LIT in CSQQDEFV and the same LIT in the SSM member identifies an IBM MQ queue manager. (For information about the **MQCONN** call, see **MQCONN** - Connect queue manager. For information about the **MQCONNX** call, see **MQCONNX** - Connect queue manager (extended).)

##### **TYPE=END**

Specifies that the table is complete. If this parameter is omitted, TYPE=ENTRY is assumed.

### Using the CSQQDEFX macro

Figure 189 on page 1301 shows the general layout of a queue manager definition table.

```

CSQQDEFX NAME=subsystem1,LIT=token1
CSQQDEFX NAME=subsystem2,LIT=token2,TYPE=DEFAULT
CSQQDEFX NAME=subsystem3,LIT=token3
...
CSQQDEFX NAME=subsystemN,LIT=tokenN
CSQQDEFX TYPE=END
END

```

Figure 189. Layout of a queue manager definition table

### Setting up the IMS trigger monitor:

You can set up an IMS batch-oriented program to monitor an IBM MQ initiation queue.

Define the application to IMS using the model CSQQTAPL in the thlqual.SCSQPROC library (see Example transaction definition for CSQQTRMN ).

Generate the PSB and ACB using the model CSQQTPSB in the thlqual.SCSQPROC library (see Example PSB definition for CSQQTRMN ).

```

* This is the application definition *
* for the IMS Trigger Monitor BMP *

APPLCTN PSB=CSQQTRMN,
PGMTYPE=BATCH,
SCHDTYP=PARALLEL

```

Figure 190. Example transaction definition for CSQQTRMN

```

PCB TYPE=TP, ALTPCB for transaction messages
MODIFY=YES, To "triggered" IMS transaction
PCBNAME=CSQQTRMN
PCB TYPE=TP, ALTPCB for diagnostic messages
MODIFY=YES, To LTERM specified or "MASTER"
PCBNAME=CSQQTRMG,
EXPRESS=YES
PSBGEN LANG=ASSEM,
PSBNAME=CSQQTRMN, Runs program CSQQTRMN
CMPAT=YES

```

Figure 191. Example PSB definition for CSQQTRMN

For further information about starting and stopping the IMS trigger monitor, see Controlling the IMS trigger monitor.

## Setting up the IMS bridge

The IBM MQ - IMS bridge is an optional component that enables IBM MQ to input and output to and from existing programs and transactions that are not IBM MQ-enabled.

This topic describes what you must do to customize the IBM MQ - IMS bridge.

### Define the XCF and OTMA parameters for IBM MQ.

This step defines the XCF group and member names for your IBM MQ system, and other OTMA parameters. IBM MQ and IMS must belong to the same XCF group. Use the OTMACON keyword of the CSQ6SYSP macro to tailor these parameters in the system parameter load module.

See Using CSQ6SYSP for more information.

### Define the XCF and OTMA parameters to IMS.

This step defines the XCF group and member names for the IMS system. IMS and IBM MQ must belong to the same XCF group.

Add the following parameters to your IMS parameter list, either in your JCL or in member DFSPBxxx in the IMS PROCLIB:

#### OTMA=Y

This starts OTMA automatically when IMS is started. (It is optional, if you specify OTMA=N you can also start OTMA by issuing the IMS command /START OTMA.)

#### GRNAME=

This parameter gives the XCF group name.

It is the same as the group name specified in the storage class definition (see the next step), and in the Group parameter of the OTMACON keyword of the CSQ6SYSP macro.

#### OTMANM=

This parameter gives the XCF member name of the IMS system.

This is the same as the member name specified in the storage class definition (see the next step).

### Tell IBM MQ the XCF group and member name of the IMS system.

This is specified by the storage class of a queue. If you want to send messages across the IBM MQ - IMS bridge you must specify this when you define the storage class for the queue. In the storage class, you must define the XCF group and the member name of the target IMS system. To do this, either use the IBM MQ operations and control panels, or use the IBM MQ commands as described in Introduction to Programmable Command Formats.

### Set up the security that you require.

The /SECURE OTMA IMS command determines the level of security to be applied to every IBM MQ queue manager that connects to IMS through OTMA. See Security considerations for using IBM MQ with IMS for more information.

**Related concepts:**

“Setting up the IMS adapter” on page 1294

To use IBM MQ within IMS requires the IBM MQ - IMS adapter (generally referred to as the IMS adapter).

**Related information:**

IBM MQ and IMS

IMS and IMS bridge applications on IBM MQ for z/OS

## Using IBM MQ with CICS

To use IBM MQ with CICS, you must configure the IBM MQ CICS adapter and, optionally, the IBM MQ CICS bridge components.

For more information about configuring the IBM MQ CICS adapter and the IBM MQ CICS bridge components, see the Configuring connections to MQ section of the CICS documentation.

**Related concepts:**

“Using IBM MQ with IMS” on page 1294

The IBM MQ -IMS adapter, and the IBM MQ - IMS bridge are the two components which allow IBM MQ to interact with IMS.

**Related reference:**

“Upgrading and applying service to Language Environment or z/OS Callable Services”

The actions you must take vary according to whether you use CALLLIBS or LINK, and your version of SMP/E.

**Related information:**

IBM MQ and CICS

## Upgrading and applying service to Language Environment or z/OS Callable Services

The actions you must take vary according to whether you use CALLLIBS or LINK, and your version of SMP/E.

The following tables show you what you need to do to IBM MQ for z/OS if you upgrade your level of, or apply service to, the following products:

- Language Environment
- z/OS Callable Services (APPC and RRS for example)

*Table 159. Service has been applied or the product has been upgraded to a new release*

Product	Action if using CALLLIBS and SMP/E V3r2 or later <b>Note:</b> You do not need to run separate jobs for Language Environment and Callable services. One job will suffice.	Action if using LINK
Language Environment	<ol style="list-style-type: none"> <li>1. Set the Boundary on your SMP/E job to the Target zone.</li> <li>2. On the SMP_CNTL card specify LINK LMODS CALLLIBS. You can also specify other parameters such as CHECK, RETRY(YES) and RC. See <i>SMP/E for z/OS: Commands</i> for further information.</li> <li>3. Run the SMP/E job.</li> </ol>	No action required provided that the SMP/E zones were set up for automatic relinking, and the CSQ8SLDQ job has been run.

Table 159. Service has been applied or the product has been upgraded to a new release (continued)

Product	Action if using CALLLIBS and SMP/E V3r2 or later <b>Note:</b> You do not need to run separate jobs for Language Environment and Callable services. One job will suffice.	Action if using LINK
Callable Services	<ol style="list-style-type: none"> <li>1. Set the Boundary on your SMP/E job to the Target zone.</li> <li>2. On the SMP_CNTL card specify LINK LMODS CALLLIBS. You can also specify other parameters such as CHECK, RETRY(YES) and RC. See <i>SMP/E for z/OS: Commands</i> for further information.</li> <li>3. Run the SMP/E job.</li> </ol>	No action required provided that the SMP/E zones were set up for automatic relinking, and the CSQ8SLDQ job has been run.

Table 160. One of the products has been updated to a new release in a new SMP/E environment and libraries

Product	Action if using CALLLIBS and SMP/E V3r2 or later <b>Note:</b> You do not need to run three separate jobs for Language Environment and Callable services. One job will suffice for both products.	Action if using LINK
Language Environment	<ol style="list-style-type: none"> <li>1. Change the DDDEFs for SCEELKED and SCEESPC to point to the new library.</li> <li>2. Set the Boundary on your SMP/E job to the Target zone.</li> <li>3. On the SMP_CNTL card specify LINK LMODS CALLLIBS. You can also specify other parameters such as CHECK, RETRY(YES) and RC. See <i>SMP/E for z/OS: Commands</i> for further information.</li> <li>4. Run the SMP/E job.</li> </ol>	<ol style="list-style-type: none"> <li>1. Delete the XZMOD subentries for the following LMOD entries in the IBM MQ for z/OS target zone: CMQXDCST, CMQXRCTL, CMQXSUPR, CSQCBE00, CSQCBE30, CSQCBP00, CSQCBP10, CSQCBR00, CSQUCVX, CSQUDLQH, CSQVXPCB, CSQVXSPT, CSQXDCST, CSQXRCTL, CSQXSUPR, CSQXTDMI, CSQXTCP, CSQXTNSV, CSQ7DRPS, IMQB23IC, IMQB23IM, IMQB23IR, IMQS23IC, IMQS23IM, IMQS23IR</li> <li>2. Set up the appropriate ZONEINDEXs between the IBM MQ zones and the Language Environment zones.</li> <li>3. Tailor CSQ8SLDQ to refer to the new zone on the FROMZONE parameter of the LINK commands. CSQ8SLDQ can be found in the SCSQINST library.</li> <li>4. Run CSQ8SLDQ.</li> </ol>
Callable services	<ol style="list-style-type: none"> <li>1. Change the DDDEF for CSSLIB to point to the new library</li> <li>2. Set the Boundary on your SMP/E job to the Target zone.</li> <li>3. On the SMP_CNTL card specify LINK LMODS CALLLIBS. You can also specify other parameters such as CHECK, RETRY(YES) and RC. See <i>SMP/E for z/OS: Commands</i> for further information.</li> <li>4. Run the SMP/E job.</li> </ol>	<ol style="list-style-type: none"> <li>1. Delete the XZMOD subentries for the following LMOD entries in the IBM MQ for z/OS target zone: CMQXRCTL, CMQXSUPR, CSQBSRV, CSQILPLM, CSQXJST, CSQXRCTL, CSQXSUPR, CSQ3AMGP, CSQ3EPX, CSQ3REPL</li> <li>2. Set up the appropriate ZONEINDEXs between the IBM MQ zones and the Callable Services zones.</li> <li>3. Tailor CSQ8SLDQ to refer to the new zone on the FROMZONE parameter of the LINK commands. CSQ8SLDQ can be found in the SCSQINST library.</li> <li>4. Run CSQ8SLDQ.</li> </ol>

For an example of a job to relink modules when using CALLLIBS, see “Running a LINK CALLLIBS job” on page 1305.



## Running a LINK CALLLIBS job

An example job to relink modules when using CALLLIBS.

The following is an example of the job to relink modules when using CALLLIBs on a SMP/E V3r2 system. You must provide a JOBCARD and the data set name of SMP/E CSI that contains IBM MQ for z/OS.

```
//*****
//* RUN LINK CALLLIBS.
//*****
//CALLLIBS EXEC PGM=GIMSMP,REGION=4096K
//SMPCSI DD DSN=your.csi
// DISP=SHR
//SYSPRINT DD SYSOUT=*
//SMPCNTL DD *
SET BDY(TZONE).
LINK LMODS CALLLIBS .
/*
```

Figure 192. Example SMP/E LINK CALLLIBS job

## Using OTMA exits in IMS

Use this topic if you want to use IMS Open Transaction Manager Access exits with IBM MQ for z/OS.

If you want to send output from an IMS transaction to IBM MQ, and that transaction did not originate in IBM MQ, you need to code one or more IMS OTMA exits.

Similarly if you want to send output to a non-OTMA destination, and the transaction did originate in IBM MQ, you also need to code one or more IMS OTMA exits.

The following exits are available in IMS to enable you to customize processing between IMS and IBM MQ:

- An OTMA pre-routing exit
- A destination resolution user (DRU) exit

### OTMA exit names

You must name the pre-routing exit DFSYPRX0. You can name the DRU exit anything, as long as it does not conflict with a module name already in IMS.

#### Specifying the destination resolution user exit name

You can use the *Druexit* parameter of the OTMACON keyword of the CSQ6SYSP macro to specify the name of the OTMA DRU exit to be run by IMS.

To simplify object identification, consider adopting a naming convention of DRU0xxxx, where xxxx is the name of your IBM MQ queue manager.

If you do not specify the name of a DRU exit in the OTMACON parameter, the default is DFSYDRU0. A sample of this module is supplied by IMS. See the *IMS/ESA® Customization Guide* for information about this.

#### Naming convention for IMS destination

You need a naming convention for the destination to which you send the output from your IMS program. This is the destination that is set in the CHNG call of your IMS application, or that is preset in the IMS PSB.

## A sample scenario for an OTMA exit

Use the following topics for an example of a pre-routing exit and a destination routing exit for IMS:

- “The pre-routing exit DFSYPRX0”
- “The destination resolution user exit” on page 1307

To simplify identification, make the OTMA destination name similar to the IBM MQ queue manager name, for example the IBM MQ queue manager name repeated. In this case, if the IBM MQ queue manager name is " VCPE ", the destination set by the CHNG call is " VCPEVCPE ".

### Related concepts:

“Using IBM MQ with IMS” on page 1294

The IBM MQ -IMS adapter, and the IBM MQ - IMS bridge are the two components which allow IBM MQ to interact with IMS.

### Related information:

IBM MQ and IMS

IMS and IMS bridge applications on IBM MQ for z/OS

## The pre-routing exit DFSYPRX0

This topic contains a sample pre-routing exit for OTMA in IMS.

You must first code a pre-routing exit DFSYPRX0. Parameters passed to this routine by IMS are documented in *IMS/ESA Customization Guide*.

This exit tests whether the message is intended for a known OTMA destination (in our example VCPEVCPE). If it is, the exit must check whether the transaction sending the message originated in OTMA. If the message originated in OTMA, it will have an OTMA header, so you should exit from DFSYPRX0 with register 15 set to zero.

- If the transaction sending the message did not originate in OTMA, you must set the client name to be a valid OTMA client. This is the XCF member-name of the IBM MQ queue manager to which you want to send the message. The *IMS/ESA Customization Guide* tells you where to set this value. We suggest you set your client name (in the OTMACON parameter of the CSQ6SYSP macro) is set to the queue manager name. This is the default. You should then exit from DFSYPRX0 setting register 15 to 4.
- If the transaction sending the message originated in OTMA, and the destination is non-OTMA, you should set register 15 to 8 and exit.
- In all other cases, you should set register 15 to zero.

If you set the OTMA client name to one that is not known to IMS, your application CHNG or ISRT call returns an A1 status code.

For an IMS system communicating with more than one IBM MQ queue manager, you should repeat the logic for each IBM MQ queue manager.

Sample assembler code is shown in Figure 193 on page 1307:

```

TITLE 'DFSYPRX0: OTMA PRE-ROUTING USER EXIT'
DFSYPRX0 CSECT
DFSYPRX0 AMODE 31
DFSYPRX0 RMODE ANY
*
SAVE (14,12),,DFSYPRX0&SYSDATE&SYSTIME
SPACE 2
LR R12,R15 MODULE ADDRESSABILITY
USING DFSYPRX0,R12
*
L R2,12(,R1) R2 -> OTMA PREROUTE PARMS
*
LA R3,48(,R2) R3 AT ORIGINAL OTMA CLIENT (IF ANY)
CLC 0(16,R3),=XL16'00' OTMA ORIG?
BNE OTMAIN YES, GO TO THAT CODE
*
NOOTMAIN DS 0H NOT OTMA INPUT
LA R5,8(,R2) R5 IS AT THE DESTINATION NAME
CLC 0(8,R5),=C'VCPEVCPE' IS IT THE OTMA UNSOLICITED DEST?
BNE EXIT0 NO, NORMAL PROCESSING
*
L R4,80(,R2) R4 AT ADDR OF OTMA CLIENT
MVC 0(16,R4),=CL16'VCPE' CLIENT OVERRIDE
B EXIT4 AND EXIT
*
OTMAIN DS 0H OTMA INPUT
LA R5,8(,R2) R5 IS AT THE DESTINATION NAME
CLC 0(8,R5),=C'VCPEVCPE' IS IT THE OTMA UNSOLICITED DEST?
BNE EXIT8 NO, NORMAL PROCESSING
*
EXIT0 DS 0H
LA R15,0 RC = 0
B BYEBYE
*
EXIT4 DS 0H
LA R15,4 RC = 4
B BYEBYE
*
EXIT8 DS 0H
LA R15,8 RC = 8
B BYEBYE
*
BYEBYE DS 0H
RETURN (14,12),,RC=(15) RETURN WITH RETURN CODE IN R15
SPACE 2
REQUATE
SPACE 2
END

```

Figure 193. OTMA pre-routing exit assembler sample

## The destination resolution user exit

This topic contains a sample destination resolution user exit for IMS.

If you have set registers 15 to 4 in DFSYPRX0, or if the source of the transaction was OTMA *and* you set Register 15 to zero, your DRU exit is invoked. In this example, the DRU exit name is DRU0VCPE.

The DRU exit checks if the destination is VCPEVCPE. If it is, it sets the OTMA user data (in the OTMA prefix) as follows:

**Offset OTMA user data  
(decimal)**

- 0 OTMA user data length (in this example, 334)
- 2 MQMD
- 326 Reply to format

These offsets are where the IBM MQ - IMS bridge expects to find this information.

We suggest that the DRU exit is as simple as possible. Therefore, in this sample, all messages originating in IMS for a particular IBM MQ queue manager are put to the same IBM MQ queue.

If the message needs to be persistent, IMS must use a synchronized transaction pipe. To do this, the DRU exit must set the OUTPUT flag. For further details, refer to the *IMS/ESA Customization Guide*.

Write an IBM MQ application to process this queue, and use information from the MQMD structure, the MQIIH structure (if present), or the user data, to route each message to its destination.

A sample assembler DRU exit is shown in Figure 194 on page 1309.

```

TITLE 'DRU0VCPE: OTMA DESTINATION RESOLUTION USER EXIT'
DRU0VCPE CSECT
DRU0VCPE AMODE 31
DRU0VCPE RMODE ANY
*
SAVE (14,12),,DRU0VCPE&SYSDATE&SYSTEMTIME
SPACE 2
LR R12,R15 MODULE ADDRESSABILITY
USING DRU0VCPE,R12
*
L R2,12(,R1) R2 -> OTMA DRU PARMS
*
L R5,88(,R2) R5 ADDR OF OTMA USERDATA
LA R6,2(,R5) R6 ADDR OF MQMD
USING MQMD,R6 AS A BASE
*
LA R4,MQMD_LENGTH+10 SET THE OTMA USERDATA LEN
STH R4,0(,R5) = LL + MQMD + 8
*
 CLEAR REST OF USERDATA
MVI 0(R6),X'00' ...NULL FIRST BYTE
MVC 1(255,R6),0(R6) ...AND PROPAGATE IT
MVC 256(MQMD_LENGTH-256+8,R6),255(R6) ...AND PROPAGATE IT
*
VCPE DS 0H
CLC 44(16,R2),=CL16'VCPE' IS DESTINATION VCPE?
BNE EXIT4 NO, THEN DEST IS NON-OTMA
MVC MQMD_REPLYTOQ,=CL48'IMS.BRIDGE.UNSOLICITED.QUEUE'
MVC MQMD_REPLYTOQMGR,=CL48'VCPE' SET QNAME AND QMGRNAME
MVC MQMD_FORMAT,MQFMT_IMS SET MQMD FORMAT NAME
MVC MQMD_LENGTH(8,R6),MQFMT_IMS_VAR_STRING
*
 SET REPLYTO FORMAT NAME
B EXIT0
*
EXIT0 DS 0H
LA R15,0 SET RC TO OTMA PROCESS
B BYEBYE AND EXIT
*
EXIT4 DS 0H
LA R15,4 SET RC TO NON-OTMA
B BYEBYE AND EXIT
*
BYEBYE DS 0H
RETURN (14,12),,RC=(15) RETURN CODE IN R15
SPACE 2
REQUATE
SPACE 2
CMQA EQUONLY=NO
CMQMDA DSECT=YES
SPACE 2
END

```

Figure 194. Sample assembler DRU exit

## Index

### Special characters

\*PUBLIC 893  
&ZSEL 1246

### A

access method services (AMS), defining  
page sets 1219  
accessing queue managers outside  
cluster 976

accounting

starting automatically 1230

ACPI (Advanced Configuration and  
Power Interface) 779

Active Directory

accessing client-connection channel  
definitions 720

- active log
    - input buffer size (INBUFF) 1233
    - number of buffers per write 1235
    - number of logs required 158
    - output buffer
      - number filled (WRTHRSH) 1235
      - size (OUTBUFF) 1234
    - placement 159
    - single or dual (TWOACTV) 1234
    - size of logs required 158
  - add routing entry 899
  - adding a new queue manager to a cluster using DHCP 928
  - adding a queue manager that hosts a queue 931
  - adding a structure 1213
  - adding features
    - using Installation Launchpad 293
  - address space storage requirements 134
  - addrtrge 899
  - administering JMS objects 1159
  - administration
    - commands 1163
    - verbs 1163
  - administration tool
    - configuration file 1160
    - configuring 1160
    - overview 1159
    - properties 1160
  - administration, channel 832
  - administrative structure 1213
  - administrator authority 219
  - administrators group 219
  - Adopting an MCA 842
  - Advanced Configuration and Power Interface (ACPI) 779
  - affinities 1001
  - age, specifying for OTMA 1228
  - AIX
    - additional settings required 239
    - client installation 319
      - silent 320
    - server installation 256
      - silent 258
    - uninstalling
      - product 389
  - ALCUNIT parameter of CSQ6ARVP 1236
  - aliases
    - examples of use 975
    - queue 987
    - queue-manager 984
    - reply-to queue 987
  - AllQueueManagers stanza, mqs.ini 765, 775
  - alter channel
    - z/OS 1269
  - ALTER QMGR command 1021
  - alternative-site recovery 176
  - AMQCLCHL.TAB 716
  - AMS (access method services), defining
    - page sets 1219
  - AnyNet 877
  - APF authorization of load libraries 1202
  - API call
    - identifying 724
  - ApiExitCommon stanza, mqs.ini 780
  - ApiExitLocal stanza, qm.ini 796
  - ApiExitTemplate stanza, mqs.ini 780
  - APPC
    - applying service 1303
  - APPC/MVS, defining a connection 1288, 1294
  - application programming
    - naming conventions for queues 411
  - applications using clusters 1001
  - applications, trusted 854
  - applying maintenance 646, 649
  - archive initialization parameters, setting 1222
  - archive log
    - archiving to DASD 161
    - archiving to tape 160
    - cataloging (CATALOG) 1238
    - compacting (COMPACT) 1238
    - data set
      - name prefix 1236
      - password protection (PROTECT) 1240
      - time stamp (TSTAMP) 1240
    - deallocate period 1233
    - device type (UNIT) 1241
    - input buffer size (INBUFF) 1233
    - maximum number in BSDS (MAXARCH) 1233
    - maximum number of tape units 1233, 1234
    - mounting, WTOR (ARCWTOR) 1237
    - optimize tape handling 1233
    - output buffer size (OUTBUFF) 1234
    - quiesce time (QUIESCE) 1240
    - retention period (ARCRETN) 1237
    - route codes (ARCWRTC) 1237
    - single or dual (TWOARCH) 1235
    - space allocation
      - block size (BLKSIZE) 1238
      - primary (PRIQTY) 1239
      - secondary (SECQTY) 1240
      - units (ALCUNIT) 1236
    - storage required 160
    - tape or DASD 160
    - using SMS with 161
  - archive parameter
    - default 1235
    - setting 1235
  - archiving, controlling using OFFLOAD parameter of CSQ6LOGP 1234
  - ARCPFX1 parameter of CSQ6ARVP 1236
  - ARCPFX2 parameter of CSQ6ARVP 1237
  - ARCRETN parameter of CSQ6ARVP 1237
  - ARCWRTC parameter of CSQ6ARVP 1237
  - ARCWTOR parameter of CSQ6ARVP 1237
  - assured delivery 846
  - authorities 314, 400, 403, 404, 519, 527, 632
  - authorization 299, 893
    - domain accounts 299
    - users 298
    - Windows 2003 accounts 298
  - authorization, display 893
  - auto-definition of channels 27
  - auto-definition user exit 905
  - automatic channel reconnect for TCP/IP 1262
  - automatic client reconnection 1038
  - availability, increasing 988
- ## B
- back out in-doubt messages
    - IBM i 892
    - UNIX systems 865
    - Windows systems 865
  - backing up queue manager data 1150
  - backup
    - coupling facility structures 171
    - frequency 170
    - fuzzy 172, 177
    - page sets 172
    - planning 169
    - point of recovery 170
    - tips 170
    - using DFHSM 175
    - what to back up 170
  - backup version
    - restoring 650
  - base function 1195
  - batch
    - improving application portability 1244
    - testing customization 1252
  - batch assembler, IVP 1252
  - Batch/TSO adapter
    - installing 1244
    - z/OS SNAP dump 1244
  - benefits
    - workload balancing 988
  - bind type 854
  - binding 903
  - BINDING channel state 837
  - binding, fastpath 854
  - BLKSIZE parameter of CSQ6ARVP 1238
  - bridging between clusters 20, 983
  - browsing a channel 884
  - BSDS
    - creating 1218
    - maximum number of log volumes (MAXARCH) 1233
    - preparation 1218
    - single or dual (TWOBSDS) 1235
  - buffer
    - input buffer size (INBUFF) 1233
    - number filled before write to log 1235
    - output buffer size (OUTBUFF) 1234
  - buffer manager
    - suppressing console messages 1247
  - buffer pool
    - defining 1216
  - buffer pools
    - effect on restart time 143
    - planning 143
  - byte 10
    - identifying API call in channel exit program 724

## C

- C and C++, testing customization 1261
- calculating the size of logs 1136
- callable services, applying service 1303
- CATALOG parameter of CSQ6ARVP 1238
- catalog, archive log (CATALOG) 1238
- CCSID
  - queue manager (QMCCSID) 1228
- CD-ROM
  - installing client from 319
  - installing server from 256
- CF structure
  - adding 1213
  - naming 1213
- CF structures
  - example definition statements 150
  - naming conventions 411
  - planning 147
  - recovery 173
  - size 147
  - using more than one 147
- CFRM policy
  - activating 1213
  - data set 1213
- CFSTRUCT name, what to specify 1213
- CHANGE (administration verb) 1163
- change definition, channel 887
- change log inventory utility (CSQJU003)
  - BSDS preparation 1218
  - log data set preparation 1218
- Change option 887
- changing
  - the default queue manager 692
- channel
  - administration 832, 974
  - alter
    - z/OS 1269
  - auto-definition 905
  - browsing 884
  - change definition 887
  - channel control function
    - UNIX systems 856
    - Windows systems 856
  - control commands 832
  - copy definition 888
  - create definition
    - IBM i 888
  - creating 860, 880
  - creating your own defaults 888
  - default values supplied by IBM MQ for IBM i 888
  - define
    - z/OS 1268, 1271
  - definition
    - on the client and server 713
    - using MQSC commands 713
  - delete 888
    - z/OS 1270
  - disconnect interval 974
  - display
    - IBM i 888
  - display status
    - IBM i 889
  - display, z/OS 1269
  - displaying 861, 888
- channel (*continued*)
  - displaying settings
    - IBM i 889
    - UNIX Systems 861
    - Windows systems 861
  - displaying status 889
    - IBM i 889
    - UNIX Systems 861
    - Windows systems 861
  - enabling 834
  - end 891
  - error 839
    - restarting after 844
  - fastpath binding 854
  - IBM i
    - resolve 892
  - in distributed queuing 14
  - in doubt 844
  - in-doubt channels 844
  - initiator
    - AIX, HP-UX, Solaris, and Windows systems 851
    - starting 851
    - stopping 851
    - UNIX systems, and Windows systems 851
    - z/OS 1272
  - listener
    - start, IBM i 890
    - start, z/OS 1274
    - stop, z/OS 1274
    - STRMQMLSR command 890
  - monitoring 832
  - ping
    - IBM i 890
    - UNIX systems 862
    - Windows systems 862
    - z/OS 1276
  - preparing 833
  - quiescing 842
  - renaming
    - IBM i 886
    - UNIX Systems 864
    - Windows systems 864
  - reset
    - z/OS 1277
  - Reset
    - IBM i 892
    - UNIX systems 864
    - Windows systems 864
  - resolving
    - IBM i 892
    - UNIX Systems 865
    - Windows systems 865
    - z/OS 1278
  - restart 834
  - restarting 974
  - restarting when stopped 844
  - run 862, 883
  - selecting 883
  - start 834
    - IBM i 883, 890
    - UNIX Systems 862
    - Windows systems 862
    - z/OS 1275
  - state 834, 837
  - stopping 842, 891
- channel (*continued*)
  - IBM i 891
  - UNIX systems 863
  - Windows systems 863
  - z/OS 1278
  - test, z/OS 1276
  - triggering 848
    - z/OS 1263
  - trusted 854
- channel control function 832
  - UNIX systems 856
  - Windows systems 856
- channel definition
  - on the client 714
  - on the server 716
- channel exits
  - ClientExitPath 722
  - ExitsDefaultPath 722
  - overview 721
  - path to exits 722
  - receive 721
  - security 721
  - send 721
- channel initiator
  - defining the procedure 1210
  - display, z/OS 1270
  - installation verification program 1257
  - retries 839
  - runmqchi command, IBM MQ for Windows 862
  - runmqchi command, IBM MQ on UNIX systems 862
  - start, IBM i 890
  - start, Windows systems and UNIX systems 851
  - start, z/OS 1272
  - stop, Windows systems and UNIX systems 851
  - stop, z/OS 1273
  - STRMQMCHLI command 890
  - tailoring parameters 1242
- channel listener
  - start, IBM i 890
  - start, z/OS 1274
  - stop, z/OS 1274
  - STRMQMLSR command 890
- channel states
  - BINDING 834
  - INACTIVE 839
- channel status
  - display, IBM i 888
  - display, z/OS 1280
- channels
  - Channels stanza, qm.ini 768
  - CHANNELS stanza, qm.ini 790
  - naming conventions 411
  - queue manager error log stanza, qm.ini 770
  - Channels stanza, qm.ini 768
  - CHANNELS stanza, qm.ini 790
  - checkpoint 893
  - checkpoint, number of log records (LOGLOAD) 1226
  - CHIADAPS
    - setting 1242

- CHIDISPS
  - setting 1242
- CHINIT
  - tailoring parameters 1242
- CICS
  - configuring an extended transactional client 710
  - testing customization 1261
- CICS adapter
  - testing customization 1261
- class of routing entry 901
- class of service 815
- client 649
  - installation 277
- client channel definition table
  - directory path 749
  - name of 750
  - purpose 716
  - where to find it 716
- client configuration 696
  - defining MQI channels
    - example mqs.ini file 712
  - using configuration file 727
    - CHANNELS stanza 734
    - ClientExitPath stanza 737
    - location of 728
    - LU62 stanza 738
    - MessageBuffer stanza 739
    - MQ file 727
    - NETBIOS stanza 738
    - SPX stanza 738
    - SSL stanza 741
    - TCP stanza 745
  - using environment variables 747
    - MQCCSID 748
    - MQCERTVPOL 748
    - MQCHLLIB 749
    - MQCHLTAB 750
    - MQIPADDRV 751
    - MQNAME 751
    - MQSERVER 751
    - MQSERVER, SPX default
      - socket 752
    - MQSERVER, TCP/IP default
      - port 752
    - MQSERVER, using 753
    - MQSSLCRYP 753
    - MQSSLFIPS 754
    - MQSSLKEYR 754
    - MQSSLRESET 755
    - MQSSPROXY 754
  - using IBM MQ environment variables 747
- client configuration file 727
  - location 728
  - path to exits 722
- client configuration files
  - stanzas
    - CDP 799
    - CHANNELS 734
    - ClientExitPath 737
    - LU62 738
    - MessageBuffer 739
    - NETBIOS 738
    - OCSP 799
    - SPX 738
    - SSL 741, 799
  - client configuration files (*continued*)
    - stanzas (*continued*)
      - TCP 745
      - TLS 741, 799
      - TMF and TMF/Gateway 746
- client installation 319
  - AIX 319
    - silent 320
  - client-connection channel
    - defining using IBM MQ Explorer 381
  - HP Integrity NonStop Server 321
  - HP-UX 323
    - non-interactive 324
  - IBMi, IBM MQ
    - uninstalling 403
  - IBMi, IBM MQ C 349
    - installing server and client 351
  - IBMi, IBM MQ Java 359
  - Linux 325
    - setting up the IBM MQ client 378
    - setting up the server
      - verifying using MQ Explorer 380
    - Solaris 333, 334
    - testing communications 382
    - UNIX and Linux systems
      - verifying 374
    - verifying using MQ Explorer 379
    - verifying, setting up server 375
    - Windows 337
      - modifying 347
      - modifying using MQParms 348
      - modifying using msixec 348
      - modifying using programs 347
      - verifying 374
- client-connection channels
  - client channel definition table 716
  - creating definitions on client 714
  - creating definitions on server 716
  - creating definitions, different platforms 713
  - defining client-connection channel 719, 720
  - defining server-connection channel 719
    - migrating to a later release 718
  - client-connection, defining 719
  - ClientExitPath 722
  - CLNTCONN 719
  - CLUSSDR
    - auto-definition 905
- cluster
  - designing 24
  - merging 945
  - naming conventions 19
  - organizing 18
  - overlapping 20
- cluster channels, z/OS 1282
- cluster queue 903
- cluster queue managers,
  - publish/subscribe
    - key roles 87, 90
    - other considerations 92
- clustered topics
  - publish/subscribe 82
- clusters
  - choosing transmission queue 807
- clusters (*continued*)
  - concentrating messages 812
  - description of 24
  - distribution lists 814
  - ExitProperties stanza attributes 776
  - message flow 803
  - networking considerations 819
  - passing messages 809
  - putting messages 806
  - queue sharing with 1292
  - reply-to queue 814
  - return routing 820
    - separating message flows 810
  - Clusters and queue-sharing groups 1292
  - clusters, use of
    - achieving workload balancing 988
    - components 902
      - publish/subscribe 61
      - workload balancing 988
  - CMDLEVEL 800
    - Configurable certification validation policy on Windows and UNIX platforms 618
  - CMDUSER parameter of CSQ6SYSP 1225
  - COBOL, testing customization 1261
  - coded character set identifier, queue manager (QMCCSID) 1228
  - COM+
    - errors 1074
  - command prefix string (CPF)
    - defining 1207
    - establishing 1205
      - in a sysplex environment 1208
  - command prefix strings (CPF) 411
  - commands
    - CL
      - DSPMQMVER 525
      - ENDMQM 519
      - ENDMQMCSVR 519
      - ENDSBS 400, 402, 519
      - RCDMQMIMG 400, 519, 526
      - STRSBS 315, 526, 527
      - WRKMQM 402, 526
      - WRKMQMCHL 518
      - WRKMQMCHST 519
    - DEFINE QALIAS 987
    - DEFINE QREMOTE 984
    - MQSC 713
  - commands, administration 1163
  - commit in-doubt messages
    - IBM i 892
    - UNIX systems 865
    - Windows systems 865
  - committed messages
    - IBM i 892
    - UNIX systems 865
    - Windows systems 865
  - common storage requirements 134
  - communication protocol
    - LU 6.2 696, 698
    - NetBIOS 696, 698
    - SPX 696, 698
    - TCP/IP 696, 698
  - communication type 696, 698
  - communications
    - configuring 696



- communications side object
  - IBM i 896, 897
  - z/OS 1288, 1294
- compact installation 277
- COMPACT parameter of CSQ6ARVP 1238
- compacting archive logs (COMPACT) 1238
- COMPAT option for OPMODE 678, 1227
- COMPLOG parameter of CSQ6LOGP 1232
- components
  - cluster 902
  - cluster, overview 24
  - distributed queuing network 14
- components installed 213
- components of
  - shared queuing 1290
- Components of shared queuing 1290
- compression, controlling using COMPLOG parameter of CSQ6LOGP 1232
- concentrating messages 812
- concentrators 6
- configuration
  - changing information 755
  - channel exits
    - connecting client to queue-sharing group 726
    - identifying the API call 724
    - path to exits 722
    - security exit on client connection 722
  - client-connection definitions
    - client channel definition table 716
    - creating on client 714
    - creating on different platforms 713
    - creating on server 716
    - defining client-connection channel 719, 720
    - defining server-connection channel 719
    - migrating to a later release 718
  - clients 696
    - defining MQI channels 712
  - clients, using environment variables 747
    - MQCCSID 748
    - MQCERTVPOL 748
    - MQCHLLIB 749
    - MQCHLTAB 750
    - MQIPADDRV 751
    - MQNAME 751
    - MQSERVER 751
- configuration (*continued*)
  - clients, using environment variables (*continued*)
    - MQSERVER, SPX default socket 752
    - MQSERVER, TCP/IP default port 752
    - MQSERVER, using 753
    - MQSSLCRYP 753
    - MQSSLFIPS 754
    - MQSSLKEYR 754
    - MQSSLPROXY 754
    - MQSSLRESET 755
  - clients, using IBM MQ environment variables 747
  - connecting applications 802
    - communication 802
    - message safety 846
    - triggering channels 848
  - connecting applications on IBM i
    - creating a transmission queue 829
    - preparation for 894
  - connecting applications on z/OS
    - defining IBM MQ objects 1263
    - preparation for 1261
  - connecting applications using distributed queuing 802
  - IBM i 763
    - attributes for changing 765
    - example mqs.ini file 772
    - example qm.ini file 772
    - files 763
  - IBM i queue manager
    - changing 767
    - channels stanza 768
    - log stanza 768
    - queue manager error log stanza 770
    - TCP stanza 771
  - MQ 687
  - MQ attributes 775
    - ACPI 779
    - all queue managers 775
    - API exits 780
    - default queue manager 776
    - exit properties 776
    - log defaults 777
    - queue managers 780, 781
  - queue managers
    - API exits 796
    - changing information 782
    - channels 790
    - default bind type 799
    - error logs 797
    - exit path 796
    - Exit properties stanza 800
    - installable services 783
    - logs 786
    - restricted mode 788
    - service components 785
    - SSL and TLS stanza 799
    - Subpool stanza 801
    - transmission protocols 793
    - XA resource managers 789
  - resource limit 245
  - sending messages to another queue manager 826
- configuration (*continued*)
  - defining channels 827
  - defining queues 828
  - starting the channel 832
  - server-connection definitions
    - accessing client-connection channel definitions 720
    - client channel definition table 716
    - creating on different platforms 713
    - creating on server 716
    - defining client-connection channel 719
    - defining server-connection channel 719
    - migrating to a later release 718
    - on server 713
  - setting up communications, UNIX systems 873
    - deciding a connection 873
    - LU 6.2 connection 877
    - LU 6.2 connection, receiving 878
    - LU 6.2 connection, sending 878
    - TCP connection 873
    - TCP connection, receiving 873, 1283
    - TCP connection, receiving using the listener 874
    - TCP connection, receiving using the listener backlog 875, 876
    - TCP connection, receiving using the SO\_KEEPALIVE option 877
    - TCP connection, sending 873
  - setting up communications, Windows systems 865
    - LU 6.2 connection 868
    - LU 6.2 connection, receiving 869
    - LU 6.2 connection, sending 869
    - NetBIOS connection 870
    - NetBIOS connection, defining local name 870, 871
    - NetBIOS connection, establishing LAN adapter number 871
    - NetBIOS connection, initiating 872
    - NetBIOS connection, target listener 872
    - TCP connection 866
  - UNIX systems
    - changing 756
    - editing configuration files 757
    - mqinst.ini 762
    - mqs.ini 758
    - qm.ini 759
    - using distributed queuing on IBM MQ
      - channel initiator 851
      - creating a transmission queue 829
      - triggering channels 848
      - triggering, example definitions for 848
  - configuration file 852
    - for administration tool 1160
  - configuration files 793
    - AllQueueManagers stanza, mqs.ini 765, 775
    - ApiExitCommon, mqs.ini 780
    - ApiExitLocal, qm.ini 796

- configuration files (*continued*)
  - ApiExitTemplate, mqs.ini 780
  - backing up of 692
  - Channels stanza, qm.ini 768
  - CHANNELS stanza, qm.ini 790
  - client
    - See client configuration file
  - ClusterQueueAccessControl stanza, mqs.ini 781
  - databases, qm.ini 789
  - DefaultQueueManager stanza, mqs.ini 765, 776
  - editing 756, 757, 763
  - example mqs.ini file 772
  - example mqs.ini file, MQSeries for UNIX systems 758
  - example qm.ini file 773
  - ExitPath stanza, qm.ini 796
  - ExitProperties stanza, mqs.ini 776
  - Log stanza, qm.ini 768, 786
  - LogDefaults stanza, mqs.ini 777
  - LU62 stanza, qm.ini 793
  - mqs.ini, description of 758
  - NETBIOS stanza, qm.ini 793
  - preconnect 727
  - priorities 757
  - queue manager configuration file, qm.ini 759
  - queue manager error log stanza, qm.ini 770
  - QueueManager stanza, mqs.ini 780
  - RestrictedMode stanza, qm.ini 788
  - Service stanza, qm.ini 783
  - ServiceComponent stanza, qm.ini 785
  - TCP stanza, qm.ini 771, 793
  - XAResourceManager stanza, qm.ini 789
- configuring
  - logs 786
  - the administration tool 1160
  - unsupported
    - InitialContextFactory 1160
- Configuring a publish/subscribe cluster 1009
- configuring an extended transactional client
  - CICS 710
  - Microsoft Transaction Server 711
  - Tuxedo 711
  - XA compliant transaction managers 701
- configuring channel exits
  - connecting client to queue-sharing group 726
  - identifying the API call 724
  - path to 722
  - security exit on client connection 722
- configuring communications 696
- configuring IBM MQ 295
- configuring IBM MQ accounts 298
- configuring logs 768
- connecting applications 802
  - communication 802
  - IBM i
    - creating a transmission queue 829
    - preparation for 894
- connecting applications (*continued*)
  - message safety 846
  - sending messages to another queue manager 826
    - defining channels 827
    - defining queues 828
    - starting the channel 832
  - triggering channels 848
  - using distributed queuing 802
  - using distributed queuing on distributed platforms
    - channel initiator 851
    - creating a transmission queue 829
    - triggering channels 848
    - triggering, example definitions for 848
  - z/OS
    - defining IBM MQ objects 1263
    - preparation for 1261
- connecting shared queues 933
- connection
  - APPC/MVS, z/OS 1283, 1293
  - deciding upon
    - z/OS 1283, 1293
  - defining APPC/MVS (LU 6.2) 1288, 1294
  - defining LU 6.2
    - IBM i 896
    - Windows systems 868
  - defining LU 6.2 UNIX systems 877
  - LU 6.2
    - Windows 865
    - z/OS 1283, 1293
  - NetBIOS 700
    - Windows 865
  - TCP
    - IBM i 894
    - Windows 865
    - z/OS 1283, 1284, 1293
  - to queue-sharing group 726
- connection parameters, setting 1222
- CONNSWAP parameter of CSQ6SYSP 1225
- console message monitoring 1044
- control commands
  - changing the default queue manager 692
  - creating a default queue manager 688
  - creating a queue manager 688
  - crtmqm, creating a default queue manager 688
  - endmqm, stopping a queue manager 694
  - immediate shutdown 694
  - migrate broker function (strmqbrk) 676
  - preemptive shutdown 694
  - quiesced shutdown 694
  - restarting a queue manager, strmqm 695
  - starting a queue manager 693
  - stopping a queue manager, endmqm 694
  - strmqm, restarting a queue manager 695
- control commands (*continued*)
  - strmqm, starting a queue manager, 693
- control commands, channel 832
- controlled shutdown of a queue manager 694
- conversion of data 826
- coordination with adjacent systems 811
- COPY (administration verb) 1163
- Copy option 888
- coupling facility (CF)
  - adding a structure 1213
  - backup and recovery 171
  - customization 1213
  - planning the environment 146
  - structure size 147
  - testing 1256
- coupling facility structures
  - example definition statements 150
  - naming conventions 411
  - planning 147
  - size 147
  - using more than one 147
- Coupling Facility structures
  - recovery 173
- CPF (command prefix string) 411
  - defining 1207
  - establishing 1205
  - in a sysplex environment 1208
- Create option 888
- creating
  - a default queue manager 688
  - a queue manager 688
  - channel
    - IBM i 880
    - UNIX systems 860
    - Windows systems 860
  - defaults 888
  - JMS objects 1167
  - objects
    - IBM i 880
    - UNIX systems 859
    - Windows systems 859
  - procedures for IBM MQ AMS 1248
  - queues 829
    - transmission queue 829
- creating domain accounts 299
- CRTCSI command 897
- CRTMQM command 859
- CSQ\_ADMIN structure 1213
- CSQ4INSX sample program 903
- CSQ4IVP1 installation verification program
  - overview 1252
  - RACF commands 1253
- CSQ4IVPX installation verification program
  - example output 1259
  - overview 1257
  - RACF commands 1258
- CSQ4MSRR sample started task procedure 1209
- CSQ4MSTR sample started task procedure 1209
- CSQ4PAGE page set sample 1219
- CSQ5PQSG (queue-sharing group utility) 1220

- CSQ6ARVP macro 1221, 1235
- CSQ6LOGP macro 1221, 1232
- CSQ6SYSP
  - macro 1221, 1223
  - QINDEXBLD parameter 1228
- CSQBDEFV 1244
- CSQINP1
  - issuing commands from 1215
  - sample data set 1215
- CSQINP2
  - issuing commands from 1215
  - using 1215
- CSQINPV
  - sample data set 1215
- CSQINPX
  - using 1215
- CSQOREXX 1245
- CSQQDEFV, subsystem definition
  - table 1295, 1300
- CSQQDEFX macro 1300
- CSQQTAPL 1301
- CSQQTPSB 1301
- CSQQTRMN 1301
- CSQSNAP 1244
- CSQWDMP 1246
- CSQWDPRD 1246
- CSQZPARM
  - creating 1221
  - displaying settings 1221
- customization 413
- customizing
  - before you start 1198
  - coupling facility related sysplex tasks 1213
  - Db2 related sysplex tasks 1211
  - IMS bridge 1302
  - initialization input data sets 1215
  - introduction 1198
  - IPL 1198
  - summary 1200
  - tasks 1198
  - testing 1252

## D

- data
  - negotiation 832
- data compression 721
- data conversion 853
  - ConvEBCDICNewline attribute, AllQueueManagers stanza 765, 775
  - EBCDIC NL character conversion to ASCII 765, 775
- data encryption 721
- data partitioning 931
- data sets
  - bootstrap, creating 1218
  - initialization 1215
  - log, creating 1218
- data-sharing group, adding IBM MQ
  - entries 1220
- database (Db2)
  - drop 1211
- DB2
  - adding IBM MQ entries to the data-sharing group 1220
  - customization 1211

- DB2 (*continued*)
  - group attachment 1229
  - naming conventions 410
  - number of BLOB tasks 1229
  - number of servers 1229
  - planning the environment 155
  - RRS Attach 155
  - storage requirements 155
  - testing 1256
- dead-letter queue 820
  - UNIX systems 854
  - Windows systems 854
  - z/OS 854
- dead-letter queues
  - specifying 689
- DEALLCT parameter of CSQ6LOGP 1233
- default
  - archive parameters 1235
  - CCSID 1228
  - CSQ6ARVP macro 1235
  - CSQ6LOGP macro 1232
  - CSQ6SYSP macro 1223
  - logging parameters 1232
  - routing code 1229
  - user ID 1225
- default channel values
  - IBM i 888
- default configuration
  - Getting Started 303
- default logging
  - installation 297
- default object creation 859
- DefaultQueueManager stanza, mqs.ini 765, 776
- defaults
  - changing the default queue manager 692
  - creating a default queue manager 688
  - queue manager 689
  - reverting to the original default queue manager 692
  - transmission queue 689
- DEFINE (administration verb) 1163
- define channel
  - z/OS 1268, 1271
- DEFINE QREMOTE command 1021
- defining
  - an LU 6.2 connection
    - IBM i 896
    - UNIX systems 877
    - Windows systems 868
  - APPC/MVS (LU 6.2) connection
    - z/OS 1288, 1294
  - MQ to IMS 1296
  - objects
    - z/OS 1263
  - queues
    - z/OS 1263
  - subsystems 1205
  - z/OS 1268, 1271
- defining a client-connection channel
  - on the server 719
- defining a server-connection channel
  - on the server 719

- defining channels
  - on the client 714
  - on the client and server 713
  - on the server 716
  - using MQSC commands 713
- DELETE (administration verb) 1163
- delete channel
  - IBM i 888
  - z/OS 1270
- deletion
  - entire 402
  - Java client 403
  - standard 400
- delivery, messages 846
- deliveryIBM MQ for z/OS 413
- designing clusters 24
- destination queue 810
- destination resolution exit
  - sample 1307
  - specifying name 1227
  - writing 1305
- destination resolution user 1305
- device type for logs (UNIT) 1241
- DFHSM 175
- DFSYDRU0 sample module 1305
- DFSYPRX0 1305
- DHCP 928
- disabled receiver channels 862, 890
- disaster recovery 176
- disconnect interval 974
- Disk space requirements 99
- display
  - option 888
  - system settings 1221
  - z/OS, DQM 1270
- DISPLAY (administration verb) 1163
- display channel
  - IBM i 888
  - UNIX systems 861
  - Windows systems 861
  - z/OS 1269
- display channel initiator
  - z/OS 1270
- display channel status
  - z/OS 1280
- Display channel status
  - UNIX systems 861
  - Windows systems 861
- display DQM 1270
- DISPLAY PUBSUB command 1021
- distributed queuing 14
  - defining the data sets 1210
  - installation verification
    - program 1257
    - LE runtime library 1204
    - SCEERUN 1204
    - security 1214
    - setting the CCSID 1228
    - supplied samples 1215
    - testing customization 1257
- Distributed queuing and clusters 850
- distribution libraries, storage requirements 133
- distribution lists 814, 826
- diverting message flows 813
- DLTLICPGM command 401, 402
- domain account 219, 297

- domain account (*continued*)
  - configuring 298
- DQM
  - display, z/OS 1270
- DRU 1305
- DRU exit
  - sample 1307
  - specifying name 1227
  - writing 1305
- Druexit 1305
- druexit name, specifying for
  - OTMA 1227
- DSPMQAUT 893
- DSPMQMVER 525
- dual BSDS (TWOBSDS) 1235
- dual logging
  - specifying for active log (TWOACTV) 1234
  - specifying for archive log (TWOARCH) 1235
- dump
  - dumping log records (dmpmqlog command) 1146
  - dumping the contents of a recovery log 1146
- dump formatting member 1246
- dynamic calls, IMS 1295
- dynamic registration 709

## E

- early code
  - library 1203
- EBCDIC NL character conversion to
  - ASCII 765, 775
- edit
  - change
    - IBM i 887
  - copy
    - IBM i 888
  - create
    - IBM i 888
  - delete
    - IBM i 888
- enabling a channel to transmit
  - messages 834
- encryption
  - parameter file 292
- end 863
- END (administration verb) 1163
- End option 891
- ending
  - a queue manager 694
- ending a channel 863, 891
- ENDMQM 519
- ENDMQMCSVR 519
- ENDSBS 400, 402, 519
- environment variables 747
  - change setting 747
  - display current setting 747
  - MQCCSID 748
  - MQCHLLIB 749
  - MQCHLTAB 750
  - MQIPADDRV 751
  - MQNAME 751
  - MQSERVER 751
  - MQSPREFIX 765, 775

- environment variables (*continued*)
  - MQSSLCRYP 753
  - MQSSLFIPS 754
  - MQSSLKEYR 754
  - MQSSLPROXY 754
  - MQSSLRESET 755
- error
  - channel 839
  - conditions when creating an
    - object 1166
  - conditions when using an
    - object 1166
  - logs 863
- example
  - alias walk-through 818
  - channel names 5
  - choosing the transmission queue 807
  - communication in IBM MQ for IBM i,
    - TCP connection 894
  - concentrating messages 812
  - create channel 860
  - creating reply-to aliases 804
  - defining channels 827
  - defining queues 828
  - defining remote queue
    - definitions 804
  - display channel 861
  - display channel status 861
  - diverting message flows 813
  - flow control 803
  - multiple subscriptions 94
  - output from CSQ4IVPX 1259
  - passing messages through
    - system 809
  - propagation of publications 94
  - propagation of subscriptions 94
  - putting messages on remote
    - queues 806
  - QM-concentrators 6
  - queue name resolution 821
  - receiving messages 808
  - renaming a channel 864
  - reply-to queue 814, 816
  - sending messages 826
  - separating message flows 810
  - setting up communication for
    - Windows systems
      - defining a NetBIOS
        - connection 870
      - defining a TCP connection 866
      - defining an LU 6.2
        - connection 868
  - setting up communication in IBM MQ
    - for z/OS
      - LU 6.2 connection 1288, 1294
      - TCP connection 1284
  - setting up communication in UNIX
    - systems
      - defining an LU 6.2
        - connection 877
  - SMP/E LINK CALLLIBS job 1305
  - starting a channel 862, 883
  - triggering 849
  - using the remote queue definition
    - object 804

- examples
  - adding a new queue manager to a
    - cluster using DHCP 928
  - adding a queue manager that hosts a
    - queue 931
  - cluster topologies 16
  - communicating with other
    - clusters 983
  - communicating with queue managers
    - outside the cluster 976
  - Configuring a publish/subscribe
    - cluster 1009
  - connecting shared queues 933
  - mqs.ini file 772
  - mqs.ini file, MQSeries for UNIX
    - systems 758
  - qm.ini file 773
  - qm.ini file, IBM MQ for UNIX
    - systems 761
  - removing a queue from a queue
    - manager 965
- EXCLMSG parameter of
  - CSQ6SYSP 1225
- exit program
  - number of TCBS 1226
  - time allowed per invocation 1226
- EXITLIM parameter of CSQ6SYSP 1226
- ExitPath stanza, qm.ini 796
- ExitProperties stanza, mqs.ini 776
- exits
  - overview 721
  - receive 721
  - security 721
    - client connection 722
  - send 721
- ExitsDefaultPath 722
- EXITTCB parameter of CSQ6SYSP 1226
- extended transactional client
  - configuring
    - CICS 710
    - Microsoft Transaction Server 711
    - Tuxedo 711
    - XA compliant transaction
      - managers 701
- EXTSHM, using 128

## F

- fast, nonpersistent messages 846
- sequence of retrieval 822
- features
  - adding
    - using Installation Launchpad 293
  - for a server installation 215
  - removing
    - using Installation Launchpad 293
- file sizes, for logs 1136
- files
  - logs 1130
  - MQ configuration 758
  - queue manager configuration 759
  - sizes, for logs 1136
- finding message affinities 1001
- first failure support technology
  - (FFST) 240
- flow control 803

- for shared queuing
  - listeners 1290
- formatting dumps 1246
- French language letter 1195
- French language support 409
- function keys, updating 1245
- fuzzy backup 172, 177

## G

- gateway between clusters 20, 983
- Getting Started help 303
- global trace
  - initial setting 1231
  - start automatically 1230
- granting
  - RACDCERT permissions to security administrator 1251
  - users resource permissions 1251
- group name, specifying for OTMA 1227
- group profiles
  - QMOMADM 314, 402, 526, 527
- groups
  - creating 219
- GRTMQUAUT 893
- gsk7capicmd.exe 603
- gsk7cmd.exe 603
- gsk7ikm.exe 603
- guidelines for creating queue managers 688

## H

- handling message affinities 1001
- Help Center 303
- High availability 1037
- high availability cluster
  - active node 1048
  - failover 1048
  - standby 1048
  - standby node 1048
  - take over 1048
- HP Integrity NonStop Server
  - client installation 321
  - file system 221
- HP Integrity NonStop systems
  - hardware and software requirements 231
- HP-UX
  - additional settings required 240
  - client installation 323
    - non-interactive 324
  - server installation 259
    - non-interactive 261
  - uninstalling
    - product 391

## I

- IBM i
  - configuration 763
    - attributes for changing 765
    - example mqs.ini file 772
    - example qm.ini file 772
    - files 763

- IBM i (*continued*)
  - connecting applications
    - creating a transmission queue 829
    - preparation for 894
  - queue manager configuration
    - channels stanza 768
    - log stanza 768
    - queue manager error log stanza 770
    - TCP stanza 771
- IBM MQ Explorer
  - %NOREPOS% 18
- IBM MQ for z/OS
  - reset channel sequence numbers 1277
  - resolving in-doubt message on channel 1278
- IBM MQ for z/OS Unix System Services Components 1195
- IBM MQ Migration Guide 419
- IBM MQ z/OS Migration Guide 419
- IBMi
  - additional settings required 247
  - installation
    - MQ components 213
  - MQ client installation 349
    - installing server and client 351
    - uninstalling 403
  - MQ Java installation 359
    - reinstalling 404
  - server installation 303, 310
    - post installation tasks 314
  - uninstalling 400
    - completely 402
    - Java 403
    - product, retaining data 400
- IBMi systems
  - hardware and software requirements 229
- IEFSSNss SYS1.PARMLIB member 1205
- importance
  - workload management 1211
- IMS
  - dynamic call stub, linking 1295
- IMS adapter 1301
  - CSQQDEFV, subsystem definition table 1300
  - CSQQDEFX, macro 1300
  - defining IBM MQ to it 1300
  - installing 1294
  - language interface token (LIT) 1300
  - SSM EXEC parameter 1296
  - SSM specification options 1299
  - subsystem member entry in IMS.PROCLIB 1296
- IMS bridge
  - age, specifying for OTMA 1228
  - customizing 1302
  - druexit name, specifying for OTMA 1227
  - group name, specifying for OTMA 1227
  - member name, specifying for OTMA 1227
  - OTMA parameters 1227
  - persistent messages 1308
  - storage class 1302

- IMS bridge (*continued*)
  - suppressing console messages 1247
  - Tpipe name 1228
- IMS trigger monitor 1301
- IMS.PROCLIB library 1296
- in-doubt channels, manual resynchronization 844
- in-doubt message on channel, resolve on z/OS 1278
- in-doubt messages, commit or back out
  - IBM i 892
  - UNIX systems 865
  - Windows systems 865
- INACTIVE channel state 837, 839
- inbound channels
  - with shared queuing 1291
- Inbound channels with shared queuing 1291
- INBUFF parameter of CSQ6LOGP 1233
- increasing availability 988
- information messages, suppressing 1247
- initial data negotiation 832
- INITIAL\_CONTEXT\_FACTORY
  - property 1160
- initialization file 126, 852
- initialization input data sets
  - customizing 1215
  - formats 1215
- initiator for channel
  - AIX, HP-UX, Solaris, and Windows systems 851
  - UNIX systems and Windows systems 851
  - z/OS 1272
- input buffer size (INBUFF) 1233
- installation
  - AIX
    - MQ components 195
  - AIX server 256
    - silent 258
  - Batch/TSO adapter 1244
  - bootstrap data set 1218
  - client 277, 319
    - AIX 319
    - AIX, silent 320
    - HP Integrity NonStop Server 321
    - HP-UX 323
    - HP-UX, non-interactive 324
    - Linux 325
    - Solaris 333, 334
    - verifying 374
    - verifying, setting up server 375
    - Windows 337
    - Windows, modifying 347
    - Windows, modifying using MQParms 348
    - Windows, modifying using msisexec 348
    - Windows, modifying using programs 347
  - compact 277
  - components 213
  - creating the logging environment 1221
  - CSQ6ARVP macro 1235
  - CSQ6LOGP macro 1232
  - CSQ6SYSP macro 1223

- installation (*continued*)
  - custom 277
    - client 319
    - server 256
  - defining page sets 1219
  - easy
    - client 319
    - server 256
  - from CD-ROM
    - client 319
    - server 256
  - HP-UX
    - IBM MQ components 198
  - HP-UX server 259
    - non-interactive 261
  - IBM MQ MQI client
    - IBMi 349
  - IBMi
    - MQ client 349
    - MQ client, installing with server 351
    - MQ client, uninstalling 403
  - IBMi MFT 310
  - IBMi server 303
    - post installation tasks 314
  - IMS adapter 1294
  - Linux 267, 328
    - MQ components 200
  - Linux server 263
  - log data set 1218
  - log file 297
  - MFT
    - IBMi 310
  - modifying
    - using Add/Remove Programs 347
    - using Installation Launchpad 293
    - using the Client DVD 347
  - MQ client
    - IBMi, installing with server 351
    - IBMi, uninstalling 403
  - MQ Java
    - IBMi 359
  - MQParms command 287
  - national languages 524
  - of maintenance packages 646, 649
  - on AIX 353
  - on HP-UX 354
  - on Linux 354
  - operations and control panels 1245
  - preparation 227
  - procedure 303
  - reinstallation 404
  - reinstalling, IBMi 404
  - resource limit configuration 245
  - RSTLPCPGM command 303
  - server 254
    - AIX 256
    - AIX, from CD-ROM 256
    - AIX, silent 258
    - AIX, System Management Interface Tool (SMIT) 256
    - compact 277
    - custom 277
    - HP-UX 259
    - HP-UX, non-interactive 261
    - IBMi 303

- installation (*continued*)
  - server (*continued*)
    - IBMi, post installation tasks 314
    - IBMi, reinstalling 404
    - Linux 263
    - Solaris 272
    - Solaris, silently 275
    - typical 277
    - UNIX and Linux systems, creating the user ID 234
    - Windows 277
    - Windows, creating a response file with msiexec 287
    - Windows, encrypting parameter file 292
    - Windows, modifying an installation 293
    - Windows, modifying an installation using launchpad 293
    - Windows, modifying an installation using msiexec 294
    - Windows, parameter file 289
    - Windows, using MQParms 287
    - Windows, using msiexec 279
    - Windows, using the launchpad 215
  - server types 215
  - Solaris
    - MQ components 203
  - Solaris server 272
    - silently 275
  - typical 277
  - uninstalling 389
  - uninstalling AIX 389
  - uninstalling HP-UX 391
  - uninstalling IBMi 400
    - completely 402
    - Java 403
    - retaining data 400
  - uninstalling Linux 392
  - uninstalling Solaris 394
  - uninstalling Windows 395
  - uninstalling Windows server
    - using installation process 399
    - using MQParms 399
    - using msiexec 397
    - using programs 397
  - UNIX and Linux system
    - what to install 194
  - UNIX and Linux systems
    - creating the user ID 234
    - national language, displaying messages in 317
  - updating the ISPF menu 1245
  - using CSQ6LOGP macro 1232
  - verifying 525
  - Windows server 277
    - creating a response file with msiexec 287
    - modifying an installation 293
    - modifying an installation using launchpad 293
    - modifying an installation using msiexec 294
    - parameter file 289
    - parameter file, encrypting 292

- installation (*continued*)
  - Windows server (*continued*)
    - using MQParms 287
    - using msiexec 279
    - using the launchpad 215
  - Windows systems
    - MQ components 207
    - security considerations 219
    - what to install 194
  - z/OS 409
    - customizing your installation 413
    - delivery media 413
    - sub-capacitylicense charges 414
    - user macros 413
  - Installation
    - checking requirements 224
    - finding the latest information 225
    - miscellaneous 249
  - installation location 192
  - installation name 182
  - installation verification program (IVP)
    - distributed queuing 1257
    - queue manager 1252
    - using 413
  - installation verification, client
    - setting up client 378
    - setting up server 375
    - testing communication 382
    - UNIX, Linux, and Windows systems 374
  - installation verification, server 363
    - command line 364
    - local installation 364
    - Postcard application 370
    - Postcard application, local installation 370
    - Postcard application, server-to-server 372
    - server-to-server 367
  - installed versions display 525
  - installing
    - IMS adapter 1294
    - on IBM i 356
  - integrated file system 893
  - integrity of delivery 846
  - intercommunication
    - concepts 850
  - interprocess communication
    - resources 128
  - INTERVAL attribute of ALTER SECURITY 1216
  - IPC resources
    - EXTSHM 128
    - shared memory on AIX 128
  - IPCS job, formatting dumps 1246
  - IPCS list, updating 1246
  - IPCS VERBEXIT 1246
  - ISPF
    - installing panels permanently 1245
    - menu, updating 1245
    - operations and control panels, setting up 1245
  - ISPLLIB concatenation 1245
  - ISPMLIB concatenation 1245
  - ISPLLIB concatenation 1245
  - ISPTLIB concatenation 1245
  - ITLM 233

IVP (installation verification program)  
distributed queuing 1257  
queue manager 1252  
using 413  
IXCMIAPU, sample statements 1213

## J

Japanese language letter 1195  
Japanese language support 409  
Java client  
deleting 403  
migrating from MA88 476  
JMS  
objects, administering 1159  
objects, creating 1167  
JMSAdmin configuration file 1160  
JMSAdmin utility 1159  
JMSAdmin.config file 1159  
JNDI  
security considerations 1160  
jobs 519, 525  
journal receivers 893  
journals 893

## K

KEEPALIVE 840  
UNIX systems 877

## L

language  
msiexec installation 286, 342  
language considerations  
national language, displaying  
messages in 317  
Language Environment, applying  
service 1303  
language installation  
MQParms command 290, 346  
language interface token (LIT) 1297  
language interface token, LIT 1300  
language letter 1195  
language, national 1199  
LDAP naming considerations 1165  
libraries  
QMQM 213, 303, 400, 402, 525  
QMQMJAVA 303, 403, 476  
QMQMSAMP 213, 303, 525  
libraries, after installation 1196  
Licenses 233  
LINK CALLLIBS 1305  
link list, updating 1203  
links, wide-band 6  
Linux  
client installation 325  
server installation 263  
Ubuntu installation 267, 328  
uninstalling  
product 392  
Linux systems  
additional settings required 242  
list cluster channels, z/OS 1282  
listener 696  
listener, trusted 854

listeners  
for shared queuing 1290  
Listeners 1290  
listening on LU 6.2  
UNIX systems 878  
Windows systems 869  
z/OS 1288, 1294  
listening on NetBIOS  
Windows systems 872  
listening on TCP  
IBM i 895  
z/OS 1286  
LIT (language interface token) 1297  
LIT, language interface token 1300  
load libraries, APF authorization of 1202  
local administration  
creating a queue manager 688  
log  
error 863  
log data sets  
creating 1218  
single or dual 1232  
log file  
installation 297  
log initialization parameters,  
setting 1222  
log records, number between  
checkpoints 1226  
Log stanza, qm.ini 768, 786  
LogDefaults stanza, mqs.ini 777  
logging  
active log placement 159  
archives on tape or DASD 160  
archiving to DASD 161  
archiving to tape 160  
calculating the size of logs 1136  
checkpoints 1133  
contents of logs 1130  
locations for log files 1143  
log file reuse 1133  
media recovery 1144  
number of active log data sets 158  
parameters 689  
planning archive storage 160  
planning the environment 157  
single or dual? 157  
size of active log data sets 158  
types of 1131  
using SMS with archives 161  
what happens when a disk fills  
up? 1141  
logging parameters  
default 1232  
setting 1232  
logging, single and dual 1232  
LOGLOAD parameter of  
CSQ6SYSP 1226  
logs  
calculating the size of logs 1136  
checkpoints 1133  
configuring 768, 786  
dumping log records (dmpmqlog  
command) 1146  
dumping the contents of 1146  
format of a log 1130  
Log stanza, qm.ini 768, 786  
logging parameters 689

logs (continued)  
managing 1140, 1141  
media recovery, linear logs 1144  
number of buffers per write 1235  
output from the dmpmqlog  
command 1147  
overheads 1136  
parameters 689  
primary log extents 1130  
protecting 1146  
reuse of 1133  
secondary log extents 1130  
types of logging 1131  
types of logs 1037  
using logs for recovery 1144  
what happens when a disk fills  
up? 1141  
logs for errors 863  
loopback testing 823  
LU 6.2 696, 698  
settings  
IBM i 896  
UNIX systems 877  
Windows systems 868  
LU 6.2 connection  
setting up  
Windows 865  
z/OS 1288, 1294  
LU62 stanza, qm.ini 793

## M

MA88 SupportPac 476  
macros  
CSQ6ARVP 1221, 1235  
CSQ6LOGP 1221, 1232  
CSQ6SYSP 1221, 1223  
macros intended for customer use 413  
maintenance 646, 649  
restoring a previous version 650  
maintenance level  
querying 669  
Managed File Transfer  
installation  
z/OS 417  
planning 164  
manipulating subcontexts 1164  
MAXARCH parameter of  
CSQ6LOGP 1233  
MAXCHL  
setting 1242  
MAXCNOFF parameter of  
CSQ6LOGP 1233  
maximum  
server-connection channels 840  
MAXRTU parameter of  
CSQ6LOGP 1234  
MCA  
adopting 842  
user-written 853  
media images  
description of 1144  
recording media images 1144  
recovering damaged objects during  
start up 1144  
recovering media images 1144

- member name, specifying for
    - OTMA 1227
  - memlimit 137
  - merging clusters 945
  - message
    - committed
      - IBM i 892
      - UNIX systems 865
      - Windows systems 865
    - concentrating 812
    - diverting flows 813
    - for distribution list 814
    - passing through system 809
    - putting on remote queue 805
    - queue name translations 821
    - receiving 808
    - return routing 820
    - return to sender 847
    - routing 807
    - sending and receiving 824
    - separating flows 810
    - sequence numbering 822
    - sequential retrieval 822
    - splitting 826
    - undeliverable 847
  - message affinities
    - finding 1001
    - removing 1001
    - resolving 1001
  - message channel agent
    - user-written 853
  - message channel agents
    - with shared queuing 1291
  - Message channel agents with shared queuing 1291
  - message flow control 803
    - networking considerations 819
  - message retry 847
  - message routing code (ROUTCDE) 1229
  - messages
    - assured delivery 846
    - back out in-doubt messages
      - IBM i 892
    - commit in-doubt messages
      - IBM i 892
    - information, suppressing 1247
    - resolve in-doubt messages
      - IBM i 892
    - sending 826
    - storage requirements 140
    - suppressing 1247
  - Messages
    - back out in-doubt messages 865
    - commit in-doubt messages 865
    - resolve in-doubt messages 865
  - Microsoft Transaction Server (MTS)
    - configuring an extended transactional client 711
  - Migrating
    - MSCS 573
  - migration
    - before you start 518
    - Java client 476
    - jobs 518
  - migration, testing 1252
  - mixing clusters with non-clustered networks 975, 984
  - modifying the installation
    - using Add/Remove Programs
      - client 347
    - using Installation Launchpad 293
    - using the Client DVD 347
  - monitoring and controlling channels
    - IBM i 879
    - UNIX systems 856
    - Windows 856
    - z/OS 1264
  - monitoring channels 832
  - more than 1 instance of a queue 989
  - mounting, archive log (ARCWTOR) 1237
  - MOVE (administration verb) 1163
  - MQ
    - configuration 687
    - configuration attributes 775
      - ACPI 779
      - all queue managers 775
      - API exits 780
      - default queue manager 776
      - exit properties 776
      - log defaults 777
      - queue managers 780, 781
    - Getting Started help 303
    - Help Center 303
    - Welcome page 303
    - MQ environment variables 747
    - MQ Script (MQSC) commands 713
    - MQ server
      - name of 751
    - MQ utility program (CSQUTIL)
      - creating a client channel definition file 716
  - MQCCSID
    - what it does 748
  - MQCHLLIB
    - introduction 720
    - what is does 749
  - MQCHLTAB
    - introduction 720
    - what it does 750
  - mqclient.ini
    - See client configuration file
  - MQCNO reconnection options 1038
  - MQI API calls 1225
  - MQIBindType 854
  - MQIPADDRV
    - what it does 751
  - mqm group 219
  - MQNAME 751
  - MQOPEN call
    - specifying a destination on 1001
    - using the MQOO\_BIND\_ON\_OPEN option 1001
  - MQParms command 287
  - MQParms.ini 288, 344
  - mqqs.ini configuration file
    - AllQueueManagers stanza 765, 775
    - ApiExitCommon stanza 780
    - ApiExitTemplate 780
    - DefaultQueueManager stanza 765, 776
    - definition of 756, 763
    - editing 757
    - ExitProperties stanza 776
  - mqqs.ini configuration file (*continued*)
    - LogDefaults stanza 777
    - priorities 757
    - QueueManager stanza 780
  - mqqs.ini file 126
  - MQSC
    - commands 713
  - MQSCO structure
    - introduction 715
  - MQSERVER
    - using 714
    - what it does 751
  - MQSPREFIX, environment variable 765, 775
  - MQSSLCRYP 753
  - MQSSLFIPS 754
  - MQSSLKEYR 754
  - MQSSLPROXY
    - what it does 754
  - MQSSLRESET 755
  - MSCS
    - interaction with MSDTC 1074
    - migrating 573
    - PostOnlineCommand 1073
    - PreOfflineCommand 1073
  - MSDTC
    - interaction with MSCS 1074
  - msiexec command
    - example 280, 339
    - using a response file 281, 341
    - using a transform 286, 342
  - MTS 711
  - MULCCAPT parameter of CSQ6SYSP 1226
  - multi instance queue manager
    - active instance 1075
    - automatic client reconnection 1075
    - failover 1075
    - standby instance 1075
    - switchover 1075
    - take over 1075
  - multiple clusters 20
  - multiple installations 183
  - multiple queue definitions
    - use of 990
    - using clusters with 989
  - multiple queue managers 869
  - multiple subscriptions, example 94
  - MUSR\_MQADMIN user ID 219, 299
- ## N
- name resolution 903
    - conflicts 820
    - convention 820
    - queue name translations 821
    - restriction 815
  - NAME\_PREFIX property 1160
  - NAME\_READABILITY\_MARKER property 1160
  - namelist
    - altering 966
    - example of using 966
  - namelists
    - naming conventions 410
  - naming a CF structure 1213
  - naming considerations, LDAP 1165



- naming conventions 19, 410
- national language
  - msiexec installation 286, 342
- national language feature 1195
- national language installation
  - MQParms command 290, 346
- national language support 409
  - EBCDIC NL character conversion to ASCII 765, 775
- national language, installation 524
- NetBIOS 696, 698, 870
  - connections 700
- NetBIOS communication, limitations 294
- NetBIOS connection
  - Windows 865
- NETBIOS stanza, qm.ini 793
- network planner 6
- networking 809
- networking considerations 819
- networks 4
- NEWFUNC option for OPMODE 678, 1227
- NL character, EBCDIC conversion to ASCII 765, 775
- node centric 803
- non-clustered network
  - receiving messages from 976
  - replying to 976
  - sending messages to 984
- non-swappable 1225
- NOWAIT option on QINDXBLD 1228

## O

- object authority manager (OAM) 219
- object creation, error conditions 1166
- object use, error conditions 1166
- objects
  - creating
    - default 859
    - IBM i 880
    - UNIX systems 859
    - Windows systems 859
  - defining
    - z/OS 1263
  - description of 24
  - JMS, administering 1159
  - JMS, creating 1167
  - media images 1144
  - naming conventions 410
  - recovering damaged objects during start up 1144
  - recovering from media images 1144
  - supplied samples 1215
- objects, damaged 893
- objects, recovery 893
- OFFLOAD parameter of CSQ6LOGP 1234
- Open Transaction Manager Access 1305
- Operation messages 1225
- operations and control panels
  - changing function keys 1245
  - installing permanently 1245
  - libraries 1245
  - setting up 1245
- Operator messages 1225

- OPMODE parameter of CSQ6SYSP 678, 1227
- options
  - change 887
  - copy 888
  - create 888
  - display 888
  - display status 889
  - end 891
  - ping 890
  - reset 892
  - resolve 865
    - IBM i 892
  - start 890
- organizing a cluster 18
- OTMA
  - DRU exit sample 1307
  - pre-routing exit sample 1306
- OTMACON 1305
- OTMACON parameter of CSQ6SYSP 1227
- outbound channels
  - with shared queuing 1291
- Outbound channels with shared queuing 1291
- OUTBUFF parameter of CSQ6LOGP 1234
- output buffer, logs (OUTBUFF) 1234
- overheads, for logs 1136
- overlapping clusters 20

## P

- page sets
  - adding 1219
  - back up frequency 172
  - calculating the size 139
  - defining 1219
  - dynamic expansion 142, 1220
  - initialization input data sets 1216
  - recovery 172
  - sample 1219
- panels
  - browsing a channel
    - IBM i 884
  - changing function keys 1245
  - channel start
    - IBM i 890
  - creating a channel
    - IBM i 880
  - display
    - IBM i 888
  - display channel status 889
  - ending channel
    - IBM i 891
  - IBM i
    - Display authority 893
    - grant 893
    - record 893
    - recover 893
    - resolve 892
    - revoke 893
    - work with status 889
  - installing 1245
  - installing permanently 1245
  - libraries 1245
- panels (*continued*)
  - ping
    - IBM i 890
  - reset
    - IBM i 892
  - selecting a channel
    - IBM i 883
  - setting up 1245
  - Work with channel status
    - IBM i 886
  - work-with-channel choices
    - IBM i 887
- parameter file
  - contents 289, 345
  - creating 288, 344
  - encrypting 292
- parameters
  - MQParms command 289, 345
- parmlibs, updating 1205
- path to exits
  - client configuration file 722
- PAUSED channel state 837, 839
- performance
  - compacting archive log 1239
  - workload management 1211
- ping 890
  - UNIX systems 862
  - Windows systems 862
- ping channel
  - z/OS 1276
- ping with LU 6.2 862
- PL/I, testing customization 1261
- planning
  - alternative-site recovery 176
  - backup and recovery 169
  - buffer pools 143
  - CF structures 147
  - command prefix strings (CPF) 411
  - communications protocol 409
  - coupling facility environment 146
  - customization 413
  - Db2 environment 155
  - installation 409, 417
  - logging environment 157
  - naming conventions 410
  - national language support 409
  - SMDS environment 151
  - storage requirements 132
  - Unix System Services environment 178
  - USS 178
- plans (Db2), customization 1211
- port
  - IBM MQ for z/OS 1272
- post installation, Windows 294
  - configuring 294
  - configuring accounts 298
  - creating and setting up domain accounts 299
  - final installation tasks 297
  - information for domain administrators 299
  - prepare IBM MQ wizard 295
  - using IBM MQ remotely 295
  - using IBM MQ Welcome pages 303
  - using the default configuration wizard 302

- post installation, Windows *(continued)*
  - using the Help Center 303
- postcard 303
- PostOnlineCommand 1073
- PPT (program properties table)
  - example 1205
  - updating 1205
- pre-routing exit 1305
- preconnect stanza 741
- preemptive shutdown of a queue manager 694
- PreOfflineCommand 1073
- preparation
  - getting started
    - UNIX systems 859
    - Windows systems 859
- Prepare IBM MQ wizard
  - log file 297
  - server 295
- preparing channels 833
- primary installation 184
  - Unix links 188
  - Windows features 191
- PRIQTY parameter of CSQ6ARVP 1239
- private region storage 135
- private region storage usage 136
- problem determination
  - problems with shutdown 694
- procedures
  - channel initiator 1210
  - queue manager 1209
- procedures, creating 1248
- process definition for triggering
  - z/OS 1263
- process definitions, queue-sharing group 584
- proclibs 1209
- product libraries, storage requirements 133
- product status 669
- programming considerations
  - affinities 1001
  - general 1001
- PROTECT parameter of CSQ6ARVP 1240
- PROVIDER\_PASSWORD property 1160
- PROVIDER\_URL property 1160
- PROVIDER\_USERDN property 1160
- proxy subscriptions 63, 68, 93
- publication propagation, example 94
- publish/subscribe
  - scope 1009
  - system queue errors 97
- publish/subscribe cluster queue managers, key roles 87, 90
- publish/subscribe cluster queue managers, other considerations 92
- publish/subscribe clustered topics 82
- publish/subscribe clusters 61
- publish/subscribe clusters and hierarchies
  - proxy subscriptions 63, 68, 93
  - queue manager names 63, 68, 93
- putting messages 805
  - on remote queues 805
  - to distribution lists 814

## Q

- QINDEXBLD
  - NOWAIT option 1228
  - WAIT option 1228
- QINDEXBLD parameter of CSQ6SYSP 1228
- qm.ini configuration file 781, 793
  - ApiExitLocal stanza 796
  - Channels stanza 768
  - CHANNELS stanza 790
  - definition of 759
  - editing 757
  - ExitPath stanza 796
  - Log stanza 768, 786
  - LU62 stanza 793
  - NETBIOS stanza 793
  - priorities 757
  - queue manager error log stanza 770
  - RestrictedMode stanza 788
  - Service stanza 783
  - ServiceComponent stanza 785
  - TCP stanza 771, 793
  - XAResourceManager stanza 789
- qm.ini file 126
- QMCCSID (queue manager coded character set identifier) 1228
- QMCCSID parameter of CSQ6SYSP 1228
- QMQM library 213, 303, 400, 402, 525
- QMQM subsystem 400, 402, 519, 526, 633
- QMQM user profile 314, 400, 402, 519, 527
- QMQMADM 893
- QMQMADM group profile 314, 402, 526, 527
- QMQMJAVA library 303, 403, 476
- QMQMSAMP library 213, 303, 525
- QSGDATA (queue-sharing group data parameter) 1228
- QSGDATA parameter of CSQ6SYSP 1228
- querying the maintenance level 669
- queue
  - destination 810
  - reply-to 814
- queue manager
  - alias 804
  - receiving 808
  - aliases 984
  - coded character set identifier (QMCCSID) 1228
  - configuration changes
    - API exits 796
    - channels 790
    - default bind type 799
    - error logs 797
    - exit path 796
    - Exit properties stanza 800
    - information 782
    - installable services 783
    - restricted mode 788
    - service components 785, 786
    - SSL and TLS stanza 799
    - Subpool stanza 801
    - transmission protocols 793
    - XA resource managers 789
- queue manager *(continued)*
  - data-sharing group name 1229
  - Db2 name 1229
  - in distributed queuing 14
  - installation verification
    - program 1252
    - outside cluster 968, 970, 975
  - queue-sharing group name 1228
  - testing 1252
- queue manager alias 804
  - receiving 808
- queue manager configuration
  - changing 767
- queue manager error log stanza, qm.ini 770
- queue manager hierarchies
  - connecting 1021
- queue managers
  - activating a backup queue manager 1154
  - backing up queue manager data 1150
  - changing the default queue manager 692
  - configuration files, backing up 692
  - creating a default queue manager 688
  - creating a queue manager 688
  - default for each node 689
  - dumping the contents of a recovery log 1146
  - guidelines for creating a queue manager 688
  - immediate shutdown 694
  - limiting the numbers of 688
  - log maintenance, recovery 1130
  - naming conventions 410
  - preemptive shutdown 694
  - qm.ini files 759
  - quiesced shutdown 694
  - recording media images 1144
  - restoring queue manager data 1150, 1151
  - reverting to the original default 692
  - specifying unique names for 688
  - starting a queue manager 693
  - starting a queue manager automatically 693
  - stopping a queue manager 694
- queue name
  - translations 821
- queue-sharing
  - clusters and 1292
- queue-sharing group 20
- queue-sharing group data parameter (QSGDATA) 1228
- queue-sharing groups 726
  - adding IBM MQ entries to the data-sharing group 1220
  - bind Db2 1211
  - CF structures required 1213
  - customizing Db2 1211
  - customizing the coupling facility 1213
  - data-sharing group name 1229
  - Db2 name 1229
  - name 1228

queue-sharing groups (*continued*)  
 naming conventions 410  
 process definitions 584  
 QSGDATA parameter 1228  
 testing 1256  
 QueueManager stanza, mq.s.ini 780  
 queues  
 aliases 987  
 create a transmission queue 829  
 defaults, transmission queues 689  
 defining  
 z/OS 1263  
 more than 1 instance 989  
 naming conventions 410  
 specifying dead-letter queues 689  
 specifying undelivered-message 689  
 supplied samples 1215  
 system  
 for publish/subscribe 96  
 QUIESCE parameter of  
 CSQ6ARVP 1240  
 quiesced shutdown of a queue manager 694  
 preemptive shutdown 694  
 quiescing 519  
 quiescing channels 842

## R

RACDCERT command  
 permissions, granting 1251  
 RACF  
 commands for CSQ4IVP1 1253, 1258  
 RACF audit records written during  
 connection processing 1229  
 RCDMQMIMG 400, 519, 526  
 RCRMQMJOB 893  
 readme file 213  
 receive exit 721  
 receiving  
 messages 808, 824  
 on LU 6.2  
 UNIX systems 878  
 Windows systems 869  
 z/OS 1288, 1294  
 on TCP  
 IBM i 895  
 z/OS 1286  
 receiving messages 808, 824  
 receiving on LU 6.2  
 UNIX systems 878  
 Windows systems 869  
 z/OS 1288, 1294  
 receiving on TCP  
 IBM i 895  
 z/OS 1286  
 recovery  
 achieving specific targets 174  
 activating a backup queue manager 1154  
 alternative-site recovery 176  
 backing up IBM MQ 1151  
 backing up queue manager data 1150  
 backup frequency 170  
 CF structures 173  
 checkpoints, recovery logs 1133

recovery (*continued*)  
 CICS 175  
 Coupling Facility structures 173  
 IMS 176  
 making sure messages are not lost  
 using logs 1130  
 media images, recovering 1144  
 page sets 172  
 performance 143  
 planning 169  
 point of recovery 170  
 procedures 169  
 recovering damaged objects at other times 1144  
 recovering damaged objects during start up 1144  
 recovering from problems 1144  
 restoring queue manager data 1151  
 tips 170  
 using DFHSM 175  
 using the log for recovery 1144  
 what to back up 170  
 REFINED option for MULCCAPT 1226  
 REFRESH CLUSTER command 973  
 region error options (REO) 1297  
 region sizes 137  
 registry 852, 870  
 reinstalling 404  
 using Installation Launchpad 293  
 reinstalling, IBM i 404  
 release notes 227  
 remote queue definition 804  
 remote-queue definition  
 in distributed queuing 14  
 removing a queue from a queue manager 965  
 removing features  
 using Installation Launchpad 293  
 removing message affinities 1001  
 renaming a channel  
 IBM i 886  
 UNIX systems 864  
 Windows systems 864  
 REO (region error options) 1297  
 reply-to alias 804  
 reply-to queue 814  
 alias example 816  
 reply-to queue aliases 987  
 repository  
 lifetime of information 23  
 out of service 16  
 selecting 16  
 topologies 16  
 REQUESTING channel state 837  
 RESAUDIT parameter of  
 CSQ6SYSP 1229  
 reset 864, 892  
 reset channel sequence numbers,  
 z/OS 1277  
 resilience 1075  
 resolve in-doubt message on channel,  
 z/OS 1278  
 resolve in-doubt messages 865  
 IBM i 892  
 resolve option 865  
 IBM i 892  
 resolving message affinities 1001  
 resource group 1048  
 resource limit configuration 245  
 resource permissions, granting 1251  
 Resource Recovery Services (RRS)  
 applying service 1303  
 resources, IPC 128  
 responder  
 LU6.2 862  
 responder process 862  
 response file  
 example 281, 341  
 restart  
 performance 143  
 restarting  
 channels 834  
 restarting a queue manager 695  
 restarting stopped channels 844  
 restoring  
 for 630  
 maintenance level 630  
 restoring previous backup version 650  
 restoring queue manager data 1150  
 RestrictedMode stanza, qm.ini 788  
 retention period, archive logs  
 (ARCRETN) 1237  
 RETRY channel state 837, 839  
 return routing 820  
 return to sender 847  
 ROUTCODE parameter of  
 CSQ6SYSP 1229  
 route codes, archive log  
 (ARCWRTE) 1237  
 routing code, message (ROUTCODE) 1229  
 routing entry  
 add 900  
 class 901  
 routing entry class 901  
 routing mechanism 63, 68, 93  
 routing messages 807  
 RRS (Resource Recovery Services)  
 applying service 1303  
 RRS adapter, installing 1244  
 RSTLPCGM command 303  
 run channel 862, 883  
 run channel initiator 851  
 runmqchi command  
 AIX, HP-UX, Solaris, and Windows  
 systems 851  
 UNIX systems and Windows  
 systems 851  
 RVKMQMAUT 893

## S

SAF  
 key rings  
 task 1250  
 sample  
 CSQINP1 1215  
 CSQINP2 1215  
 CSQINPV 1215  
 CSQINPX 1215  
 data set members 1198  
 defining page sets 1219  
 destination resolution exit 1307  
 IXCMIAPU statements 1213

- sample (*continued*)
  - linking the IMS dynamic call stub 1295
  - OTMA pre-routing exit 1306
  - output from CSQ4IVPX 1259
  - SMP/E LINK CALLLIBS job 1305
  - started task procedure 1209
- sample program, CSQ4INSX 903
- SCSQxxxx contents 1196
- SECQTY parameter of CSQ6ARVP 1240
- security
  - archive log 1240
  - default user ID 1225
  - installation tasks 1214
  - installation verification program
    - distributed queuing 1257
    - queue manager 1253
  - INTERVAL attribute 1216
  - levels for exit programs 854
  - protecting log files 1146
  - restoring queue manager data 1150
  - SecurityPolicy attribute, Service stanza 783
  - TIMEOUT attribute 1216
  - UNIX and Linux systems
    - creating the user ID 234
  - Windows systems
    - installation considerations 219
- security considerations, JNDI 1160
- security exit 721
- security exits
  - client connection 722
- security policy capability 1230
- SECURITY\_AUTHENTICATION property 1160
- segmented messages 1001
- selecting a channel 883
- send exit 721
- sending
  - messages 824, 826
- sequence numbering 822
- sequence numbers
  - reset, z/OS 1277
- sequential retrieval of messages 822
- server
  - installation
    - compact 277
    - custom 277
    - typical 277
- server installation 254
  - AIX 256
    - from CD-ROM 256
    - silent 258
    - System Management Interface Tool (SMIT) 256
  - HP-UX 259
    - non-interactive 261
  - IBMi 303, 310
    - post installation tasks 314
    - reinstalling 404
  - Linux 263
  - Solaris 272
    - silently 275
  - UNIX and Linux systems
    - creating the user ID 234
  - Windows 277
- server installation (*continued*)
  - creating a response file with msiexec 287
  - modifying an installation 293
  - modifying an installation using launchpad 293
  - modifying an installation using msiexec 294
  - parameter file 289
  - parameter file, encrypting 292
  - using MQParms 287
  - using msiexec 279
  - using the launchpad 215
- server-connection channel
  - defining 719
- server-connection channels
  - client channel definition table 716
  - creating definitions on server 713, 716
  - creating definitions, different platforms 713
  - defining client-connection channel 719, 720
  - defining server-connection channel 719
  - migrating to a later release 718
- server-connection channels, maximum number 840
- service class
  - workload management 1211
- service considerations 1303
- service group 1048
- Service stanza, qm.ini 783
- service, class of 815
- ServiceComponent stanza, qm.ini 785
- setmqipw command 292
- setmqscp command
  - introduction 720
- SETSSI command 1205
- setting up
  - communication
    - IBM i 894
    - UNIX systems 873
    - Windows 865
- setting up communications
  - UNIX systems 873
    - LU 6.2 connection 877
    - LU 6.2 connection, receiving 878
    - LU 6.2 connection, sending 878
    - TCP connection, receiving 1283
  - Windows systems 865
    - LU 6.2 connection 868
    - LU 6.2 connection, receiving 869
    - LU 6.2 connection, sending 869
    - NetBIOS connection 870
    - NetBIOS connection, defining local name 870, 871
    - NetBIOS connection, initiating 872
    - NetBIOS connection, target listener 872
    - TCP connection 866
- setting up the server
  - client installation
    - verifying using MQ Explorer 380
- shared channels
  - naming conventions 411
- shared memory on AIX 128
- shared queue 20
- shared queues 726
  - adding IBM MQ entries to the data-sharing group 1220
  - CF structures required 1213
  - customizing Db2 1211
  - customizing the coupling facility 1213
  - data-sharing group name 1229
  - Db2 name 1229
  - mapping to CF structures 151
  - naming conventions 411
  - QSGDATA parameter 1228
  - queue-sharing group name 1228
  - testing 1256
- shared queuing
  - components of 1290
  - inbound channels 1291
  - message channel agents with 1291
  - outbound channels 1291
  - synchronization queue 1292
- shutting down a queue manager 694
  - a queue manager, quiesced 694
  - immediate 694
  - preemptive 694
- side object
  - IBM i 897
- simple client-connection channel
  - definition 714
- Simplified Chinese
  - language letter 1195
- Simplified Chinese language
  - support 409
- single BSDS (TWOBSDS) 1235
- single logging
  - specifying for active log (TWOACTV) 1234
  - specifying for archive log (TWOARCH) 1235
- SMDS
  - planning the environment 151
- SMF (System Management Facility)
  - CSQ6SYSP, specifying parameters 1223
  - gathering (STATIME) 1230
  - starting automatically (SMFSTAT) 1230
- SMFACCT parameter of CSQ6SYSP 1230
- SMFSTAT parameter of CSQ6SYSP 1230
- SNA communication, limitations 294
- SNAP dump, Batch/TSO adapter 1244
- SO\_KEEPALIVE
  - UNIX systems 877
- sockets
  - security 1214
- Solaris
  - client installation 333
  - non-interactive client installation 334
  - server installation 272
    - silently 275
  - uninstalling
    - product 394
- space allocation
  - archive logs, block size (BLKSIZE) 1238

- space allocation (*continued*)
    - logs, primary (PRIQTY) 1239
    - logs, secondary (SECQTY) 1240
    - units, logs (ALCUNIT) 1236
  - space requirements for maintenance 628, 640, 642
  - SPLCAP parameter of CSQ6SYSP 1230
  - splitting messages 826
  - SPX 696, 698
    - default socket 752
  - SPX communication, limitations 294
  - SSM (subsystem member)
    - entry in IMS.PROCLIB 1296
    - EXEC parameter 1296
    - specification options 1299
  - STANDARD option for MULCCAPT 1226
  - stanzas 793
    - AllQueueManagers, mqs.ini 765, 775
    - ApiExitCommon, mqs.ini 780
    - ApiExitLocal, qm.ini 796
    - ApiExitTemplate, mqs.ini 780
    - Channels, qm.ini 768
    - CHANNELS, qm.ini 790
    - ClusterQueueAccessControl, qm.ini 781
    - DefaultQueueManager, mqs.ini 765, 776
    - ExitPath, qm.ini 796
    - ExitProperties, mqs.ini 776
    - Log, qm.ini 768, 786
    - LogDefaults, mqs.ini 777
    - LU62, qm.ini 793
    - NETBIOS, qm.ini 793
    - queue manager error log, qm.ini 770
    - QueueManager, mqs.ini 780
    - RestrictedMode stanza, qm.ini 788
    - Service, qm.ini 783
    - ServiceComponent, qm.ini 785
    - TCP, qm.ini 771, 793
    - XAResourceManager, qm.ini 789
  - start
    - channel 834
      - UNIX systems 862
      - Windows systems 862
      - z/OS 1275
    - channel initiator, z/OS 1272
    - channel listener, z/OS 1274
    - option 890
  - started task
    - user 1248
  - started task procedure
    - creating for channel initiator 1210
    - creating for queue manager 1209
    - security 1209
  - started task procedure, CSQ4MSRR 1209
  - started task procedure, CSQ4MSTR 1209
  - starting
    - a queue manager 693
    - a queue manager automatically 693
  - STARTING channel state 837
  - startup procedure, CSQ4MSRR 1209
  - startup procedure, CSQ4MSTR 1209
  - state, channel 834
  - STATIME parameter of CSQ6SYSP 1230
  - statistics
    - gathering time interval 1230
  - statistics (*continued*)
    - starting automatically 1230
  - status
    - display channel 861
    - work with channel 886
  - status panels 889
  - stop
    - channel 842, 863
    - channel initiator, z/OS 1273
    - channel listener, OS/390 1274
    - channel, z/OS 1278
    - quiesce 863, 864
  - stop channel initiator 851
  - stop force 864
  - STOPPED channel state 837, 839
  - stopped channels, restarting 844
  - STOPPING channel state 837
  - storage class
    - IMS bridge 1302
  - storage classes
    - naming conventions 411
  - storage group (Db2), drop 1211
  - storage requirement, trace table 1231
  - storage requirements
    - above the bar 135
    - address space 134
    - below the bar 135
    - Channel initiator private region
      - storage usage 136
    - common 134
    - Db2 155
    - introduction 132
    - messages 140
    - product libraries 133
    - queue manager private region storage usage 135
  - strmqbrk command 676
  - STRMQM command 859
  - strmqm control command 695
  - STRSBS 315, 526, 527
  - structures, size 147
  - sub-capacity usage license charges 414
  - subcontexts, manipulating 1164
  - Subpool 801
  - subscription propagation, example 94
  - subsystem name table, updating 1205
  - SupportPac MA88 476
  - suppressing information messages 1247
  - swappable 1225
  - switch load file 710
  - switch structure 701
  - synchronization
    - with shared queuing 1292
  - Synchronization queue with shared queuing 1292
  - SYS1.PARMLIB members
    - IEFSSNss 1205
  - SYS1.PARMLIB, updating 1215
  - SYSEXEC concatenation 1245
  - sysplex
    - command prefix string (CPF) 1208
    - scope 1208
  - SYSPROC, concatenation 1245
  - system extension 854
  - system extensions
    - user-exit programs
      - UNIX systems 854
  - system extensions (*continued*)
    - user-exit programs (*continued*)
      - Windows systems 854
  - System Management Interface Tool (SMIT)
    - installing client with 319
    - installing server with 256
  - system parameters
    - displaying settings 1221
    - invoking 1221
    - setting 1223
    - tailoring 1221
    - z/OS 1201
  - system preparation 233
    - additional settings required
      - AIX 239
      - HP-UX 240
      - IBMi 247
      - Linux systems 242
    - creating file systems on UNIX and Linux systems 236
    - file system
      - HP Integrity NonStop Server 221
    - hardware and software requirements
      - HP Integrity NonStop systems 231
      - IBMi systems 229
      - UNIX and Linux 225
      - Windows systems 227
    - installation location 192
    - license requirements 233
    - Solaris resource limit
      - configuration 245
  - SYSTEM.CHANNEL.INITQ queue
    - z/OS 1264, 1265
  - SYSTEM.CHANNEL.REPLY.INFO queue 1264, 1265
- ## T
- table space (Db2), drop 1211
  - tables and index (Db2), drop 1211
  - tailoring
    - channel initiator parameters 1242
  - target libraries, storage requirements 133
  - task examples
    - adding a new queue manager to a cluster using DHCP 928
    - adding a queue manager that hosts a queue 931
    - Configuring a publish/subscribe cluster 1009
    - connecting shared queues 933
    - removing a queue from a queue manager 965
  - TCB, number for exit programs 1226
  - TCP
    - connection
      - listener backlog 875, 1286
      - listener backlog option 875, 1286
      - OpenEdition MVS sockets 1262
    - TCP connection
      - setting up
        - Windows 865
        - z/OS 1284

- TCP KEEPALIVE
  - UNIX systems 877
- TCP stanza, qm.ini 771, 793
- TCP/IP 696, 698
  - default port 752
- TCP/IP connection
  - on a client 699
- TCP/IP KEEPALIVE 840
- Terminal Services
  - Using IBM MQ 295
- test channel, z/OS 1276
- testing connections, lookback testing 823
- testing your queue manager 1252
- thlqual.SCSQxxxx, contents 1196
- time stamp, archive log (TSTAMP) 1240
- TIMEOUT security attribute 1216
- topologies 16
- Tpipe, name 1228
- TPNAME and TPPATH
  - IBM i 896
  - UNIX systems 877
  - Windows systems 868
- trace
  - starting automatically (TRACSTR) 1230
  - trace table size (TRACTBL) 1231
- tracing parameters, setting 1222
- TRACSTR parameter of CSQ6SYSP 1230
- TRACTBL parameter of CSQ6SYSP 1231
- transforms 286, 342
  - MQParms command 290, 346
- translated versions, installation 524
- transmission header
  - queue name 804
- transmission protocol
  - LU 6.2 696, 698
  - NetBIOS 696, 698
  - SPX 696, 698
  - TCP/IP 696, 698
- transmission queue
  - in distributed queuing 14
- transmission queues
  - default 689
- triggering
  - channels 848
  - z/OS 1263
- trusted applications 854
- TSO
  - formatting dumps 1246
  - improving application portability 1244
- TSTAMP parameter of CSQ6ARVP 1240
- Tuxedo
  - configuring an extended transactional client 711
- TWOACTV parameter of CSQ6LOGP 1234
- TWOARCH parameter of CSQ6LOGP 1235
- TWOBSDS parameter of CSQ6LOGP 1235
- type, bind 854
- types of logging 1131
- types of server installation 215
- typical installation 277

## U

- U.S. English (mixed case) support 409
- U.S. English (uppercase) support 409
- UAC 219
- Ubuntu installation
  - Linux 267, 328
- undeliverable message 847
- undelivered-message queue
  - UNIX systems 854
  - Windows systems 854
  - z/OS 854
- uninstalling 389
  - AIX 389
  - HP-UX 391
  - IBMi 400
    - completely 402
    - Java 403
    - retaining data 400
  - Linux 392
    - on AIX 405
    - on HP-UX 407
    - on Linux 407
  - Solaris 394
  - Windows 395
  - Windows server
    - using installation process 399
    - using MQParms 399
    - using msixec 397
    - using programs 397
- UNIT parameter of CSQ6ARVP 1241
- UNIT2 parameter of CSQ6ARVP 1241
- UNIX
  - IPC resources 128
  - KEEPALIVE 877
  - process priority 129
- UNIX and Linux systems
  - client installation
    - verifying 374
  - creating file systems 236
  - hardware and software requirements 225
  - installation
    - creating the user ID 234
    - national language, displaying messages in 317
    - what to install 194
  - resource limit configuration 245
  - Solaris
    - resource limit configuration 245
- UNIX and Linux systems installation
  - IBMi
    - MQ components 213
  - UNIX operating system
    - example mqs.ini file 758
    - example qm.ini file 761
- Unix System Services environment 178
- UNIX System Services sockets
  - security 1214
- UNIX systems
  - configuration
    - changing 756
    - editing configuration files 757
    - mqinst.ini 762
    - mqs.ini 758
    - qm.ini 759
  - setting up communications 873
    - LU 6.2 connection 877

- UNIX systems (*continued*)
  - setting up communications (*continued*)
    - LU 6.2 connection, receiving 878
    - LU 6.2 connection, sending 878
    - TCP connection, receiving 1283
- updating
  - z/OS link list 1203
  - z/OS parmlibs 1205
  - z/OS subsystem name table 1205
- updating IBM MQ 646, 649
- upgrade considerations 1303
- US English language letter 1195
- USE\_INITIAL\_DIR\_CONTEXT
  - property 1160
- user exit
  - auto-definition 905
- user ID
  - maximum age in OTMA 1228
  - MUSR\_MQADMIN 219, 299
- user ID security
  - default 1225
- user profiles
  - QMOM 314, 400, 402, 519, 527
  - QMOMADM 314, 402, 526, 527
- user-exit programs
  - security levels 854
  - system extension
    - UNIX systems 854
    - Windows systems 854
- user-written MCAs 853
- USERDATA parameter
  - z/OS 1263
- using EXTSHM 128
- Using IBM MQ remotely 295
- Using IBM MQ via Terminal Services 295
- USS 178

## V

- values supplied by IBM MQ for IBM i 888
- velocity goal
  - workload management 1211
- VERBEXIT, IPCS 1246
- verification of installation, client
  - setting up client 378
  - setting up server 375
  - testing communication 382
  - UNIX, Linux, and Windows systems 374
- verification of installation, server 363
  - command line 364
  - local installation 364
  - Postcard application 370
  - Postcard application, local installation 370
  - Postcard application, server-to-server installation 372
  - server-to-server 367
- VerificationLevel for Opmode 678, 1227
- verifying installation 525
- verifying upgrade 526

## W

- WAIT, option on QINDXBLD 1228
- WAITING channel state 837
- WebSphere MQ Migration Guide 419
- WebSphere MQ z/OS Migration Guide 419
- what to install 194
  - Windows systems 194
- which queue managers should hold repositories 16
- wide-band links 6
- Windows
  - adding a queue manager to 693
- Windows 2003 accounts 298
- Windows systems
  - client installation 337
    - modifying 347
    - modifying using MQParms 348
    - modifying using msixec 348
    - modifying using programs 347
    - verifying 374
  - hardware and software requirements 227
  - installation
    - MQ components 207
    - security considerations 219
    - what to install 194
  - server installation 277
    - creating a response file with msixec 287
    - modifying an installation 293
    - modifying an installation using launchpad 293
    - parameter file 289
    - parameter file, encrypting 292
    - using MQPARms 287
    - using msixec 279
    - using the launchpad 215
  - server post installation 294
    - configuring 294
    - configuring accounts 298
    - creating and setting up domain accounts 299
    - final installation tasks 297
    - information for domain administrators 299
    - prepare IBM MQ wizard 295
    - using IBM MQ remotely 295
    - using IBM MQ Welcome pages 303
    - using the default configuration wizard 302
    - using the Help Center 303
  - setting up communications 865
    - LU 6.2 connection 868
    - LU 6.2 connection, receiving 869
    - LU 6.2 connection, sending 869
    - NETBIOS connection 870
    - NETBIOS connection, defining
      - local name 870, 871
    - NETBIOS connection, initiating 872
    - NETBIOS connection, target listener 872
    - TCP connection 866
  - uninstalling 395

- Windows systems (*continued*)
  - uninstalling server
    - using installation process 399
    - using MQParms 399
    - using msixec 397
    - using programs 397
- WLMTIME parameter of CSQ6SYSP 1231
- WLMTIMU parameter of CSQ6SYSP 1231
- work with channel status 886
- work with status 889
- work-with-channel choices 887
- working with queue manager outside cluster 968, 970, 975, 984
- workload balancing
  - achieving 988
  - algorithm 989
  - user exit 1001
  - with multiple queue definitions 989
- workload management, service class 1211
- workload manager, queue scan interval 1231
- workload manager, queue scan interval units 1231
- writing your own message channel agents 853
- WRKCLS command 901
- WRKMQM 402, 526
- WRKMQMCHL 518
- WRKMQMCHST 519
- WRKSBSD command 900
- WRTHRSH parameter of CSQ6LOGP 1235

## X

- XA compliant transaction managers
  - configuring an extended transactional client 701
- XA switch structure
  - introduction 701
  - supplied with an extended transactional client 708
  - use by CICS 710
  - use by Tuxedo 711
- xa\_open string
  - example 704
  - format 704
  - introduction 701
  - use by CICS 710
  - use by Tuxedo 711
- XAResourceManager stanza, qm.ini 789
- XCF
  - group name, specifying for OTMA 1227
  - member name, specifying for OTMA 1227

## Z

- z/OS
  - APF authorized libraries 1202
  - connecting applications
    - defining IBM MQ objects 1263

z/OS (*continued*)

- connecting applications (*continued*)
  - preparation for 1261
- connecting client to queue-sharing group 726
- installation 409
  - customizing your installation 413
  - delivery media 413
  - Managed File Transfer 417
  - sub-capacity licence charges 414
  - user macros 413
- link list, updating 1203
- parmlibs, updating 1205
- program properties table, updating 1205
- SNAP dump 1244
- subsystem name table, updating 1205
- system parameters 1201





---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Software Interoperability Coordinator, Department 49XA  
3605 Highway 52 N  
Rochester, MN 55901  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Programming interface information

Programming interface information, if provided, is intended to help you create application software for use with this program.

This book contains information on intended programming interfaces that allow the customer to write programs to obtain the services of WebSphere MQ.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

**Important:** Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

---

## Trademarks

IBM, the IBM logo, [ibm.com](http://ibm.com)<sup>®</sup>, are trademarks of IBM Corporation, registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml). Other product and service names might be trademarks of IBM or other companies.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

This product includes software developed by the Eclipse Project (<http://www.eclipse.org/>).

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.



---

## **Sending your comments to IBM**

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Use one of the following methods to send us your comments:

- Send an email to [ibmkc@us.ibm.com](mailto:ibmkc@us.ibm.com)
- Use the form on the web here: [www.ibm.com/software/data/rcf/](http://www.ibm.com/software/data/rcf/)

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate, without incurring any obligation to you.

Include the following information:

- Your name and address
- Your email address
- Your telephone or fax number
- The publication title and order number
- The topic and page number related to your comment
- The text of your comment

IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you submit.

Thank you for your participation.





